Investigating the consequences of iterated learning in phonological typology*

Coral Hughto
University of Massachusetts, Amherst
coralwilliam@linguist.umass.edu

1 Introduction

This paper adds to a growing body of work investigating the effects of learning biases on probabilistic typological predictions in phonology. Much of this literature generates probabilistic typologies through combining a particular theory of grammar with a particular theory of learning, drawing on differences in learnability between patterns to explain disparities in frequency of attestation (Pater, 2012; Staubs, 2014; Stanton, 2016, among others). I will be building on our previous work (Hughto and Pater, 2017; Pater, 2012) which combines the Maximum Entropy grammatical framework (MaxEnt; Goldwater and Johnson (2003)), a weighted-constraint model which generates a probability distribution over competing output candidates, with an interactive, agent-based learning model in which agents learn from each other, with no target grammar.

Weighted-constraint models such as MaxEnt allow for constraint cumulativity effects (e.g. gang effect patterns), where multiple violations of (a) lower-weighted constraint(s) can "gang up" to outweigh a single violation of a higher-weighted constraint, and so are often criticized for overpredicting the range of typologically possible patterns. The influence of agent interation in learning, however, yielded an emergent bias for deterministic patterns, in which one output candidate accumulates majority probability over its competitors, as well as a bias away from cumulative constraint effects ("gang effects"). The combination of MaxEnt and interactive learning model, then, causes this weighted-constraint grammatical model to produce behavior similar to that of ranked-constraint grammars, which (typically) only select one output candidate per input, and do not allow for constraint cumulativity.

The results from the interactive learning model produced a greater match to some observed typological discrepancies than the MaxEnt grammatical model alone (Hughto and Pater, 2017), but the agents in that model have no target grammar, abstracting away from the influence of more stable, adult models in acquisition. The iterated learning model (Kirby, 2017, among others) is an agent-based learning model in which data is transmitted from "parent" to "child" across multiple generations of agents, and is designed to simulate the effect of language transmission across generations. A more realistic model would likely incorporate both language transmission and agent interaction; I leave that for future work.

The work I will present here builds on our previous findings with the interactive learning model (Hughto and Pater, 2017), adding results from simulations performed with the iterated learning model. I show that the emergent learning biases which were robust across parameter settings with the interactive learning model are present, but less robust with the iterated learning model, emerging most clearly with longer learning times between generations, and with new agents introduced into the system with constraint weights initialized at zero.

2 Iterated Learning Model

In the iterated learning model, a "child" agent learns from data produced by a "teacher" target distribu-

*Thanks to Joe Pater, Gaja Jarosz, John Kingston, and audiences at UMass and elsewhere.

Proceedings of the Society for Computation in Linguistics (SCiL) 2018, pages 182-185.
Salt Lake City, Utah, January 4-7, 2018

tion for a given number of learning steps. The child agent's grammar then becomes the target grammar for the next agent in the chain, and the process is repeated for a given number of generations.

In the simulations performed here, agents in the chain had MaxEnt grammars, equipped with a set of constraints, constraint weights, a set of inputs and corresponding outputs. Learning progressed via a gradual, error-driven learning algorithm. In each learning step, an input is sampled from a uniform probability distribution over inputs, and the teacher agent samples an output for that input according to the probability distribution defined by its MaxEnt grammar. The child agent samples an output for that same input, according to the current probability distribution defined by its own MaxEnt grammar. If the outputs don't match (i.e. if the child agent has made an error), then the child agent updates its constraint weights using stochastic gradient descent, according to the equation in Figure 1 - the learner takes the difference in violation vectors between the conflicting input-output mappings, scales it by a learning rate, then adds the scaled difference to its old constraint weights.

New Weights = Old Weights + (Teacher - Learner) * Learning Rate

Figure 1: Constraint Weight Update Rule

At the end of each simulation, the final child agent's highest probability output candidate for each input is taken as the set of winning candidates, and the distribution over languages learned over multiple runs of the simulation is taken as the predicted probabilistic typology.

3 Simulations

The iterated learning model was applied to a simple, hypothetical system with two constraints, two inputs, and two possible outputs each, shown in Figure 2.

	X	Y
A		-1
B	-1	

	X	Y
C	-1	
D		-2

Figure 2: A simple, hypothetical system

There are three possible languages in this system (labeled by the winning candidates), which are given in Table 1 along with the weighting conditions which produce them. The AC language is a gang effect - one violation of constraint Y is better than one violation of constraint X (candidate A wins), but two violations of constraint Y is worse than one violation of constraint X, (candidate C wins).

	Language	Weighting condition
i.	BC	$w(Y) > w(X)$
ii.	AD	$w(X) > 2w(Y)$
iii.	AC	$2w(Y) > w(X) > w(Y)$

Table 1: Typology of the simple, hypothetical system

For the results given below, two types of simulations were run. In one type, each child agent was initialized with constraint weights randomly sampled from a uniform distribution between 0-10. In the second type, each child agent was initialized with constraint weights of zero. Additionally, the results from two different learning step values are presented here: in one set, each child agent learned from its target grammar for 100 learning steps; in the second set, each child agent learned for 1,000 learning steps. In all sets of simulations, there were 50 generations in the chain, and 1000 runs of the simulation were performed. The initial grammars in the chains were generated by randomly sampling weights for the two constraints in the system from a uniform distribution between 0-10. Disussion of other parameter setting values is omitted here for space; these results are representative of the performance of the model.

4 Results

In the results from the iterated learning model simulations, both the initial constraint weights for new agents (random or zero) and the number of learning steps between generations of agents impacted the emergence of the bias towards deterministic grammar states, and the bias away from constraint cumulativity (here, the gang effect AC language).

For lower learning steps values (here, 100), the model results demonstrate a bias away from the gang effect AC language, but only when new child agents were initialized with constraint weights at zero (Zero-Init). When new child agents were initialized with random constraint weights (Rand-Init), there was no bias against the gang effect language.

These effects can be seen in Table 2; the distribution over languages that results from simply randomly sampling 10,000 sets of constraint weights from a uniform distribution between 0-10 (Sampling) is given as a baseline for comparison. In comparison to the baseline, the Zero-Init condition gives a lower predicted probability to the gang effect AC language, while the predicted probability given by the Rand-Init condition is barely different.

Language	Sampling	Zero-Init	Rand-Init
BC	0.5	0.71	0.67
AD	0.25	0.19	0.09
AC	0.25	0.08	0.24

Table 2: Cumulativity bias results for simulations with 100 learning steps per generation

The agents in the 100 learning step condition additionally did not demonstrate any discernable bias towards more deterministic grammar states, regardless of the agents' initial constraint weights. This can be seen in Figure 3 for the Rand-Init condition, and in Figure 4 for the Zero-Init condition. These graphs plot the average probability of the winning output candidates at the end of each generation. As the simulation progresses through the generations of agents, there is no visible trend towards increasing the probability of the winning output candidates.

For higher learning step values (here, 1,000), the model results demonstrate both a bias away from the gang effect AC language, and a bias towards more deterministic grammar states, regardless of whether new child agents were initialized with random constraint weights (Rand-Init) or constraint weights of zero (Zero-Init). The bias away from constraint cumulativity can be seen in Table 3. Both the Zero-Init and Rand-Init conditions give a lower predicted probability to the gang effect AC language, compared to the sampled baseline distribution.

Language	Sampling	Zero-Init	Rand-Init
BC	0.5	0.65	0.55
AD	0.25	0.32	0.30
AC	0.25	0.03	0.15

Table 3: Cumulativity bias results for simulations with 100 learning steps per generation

The agents in the 1,000 learning step condition additionally did demonstrate a bias towards more

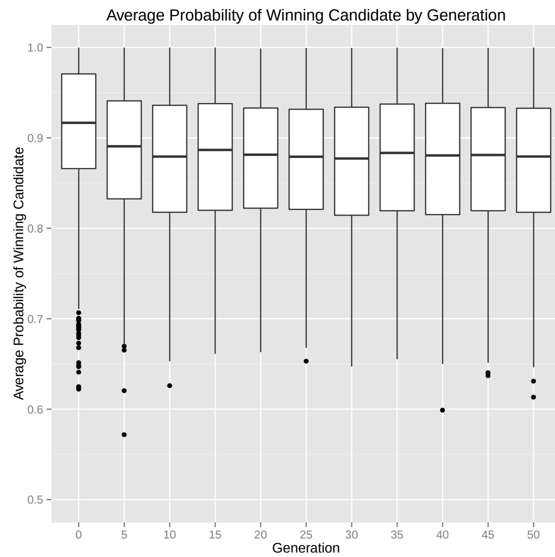

Figure 3: Average probability on the winning candidate at the end of each generation, for 100 learning steps per generation, with random initial constraint weights for new agents

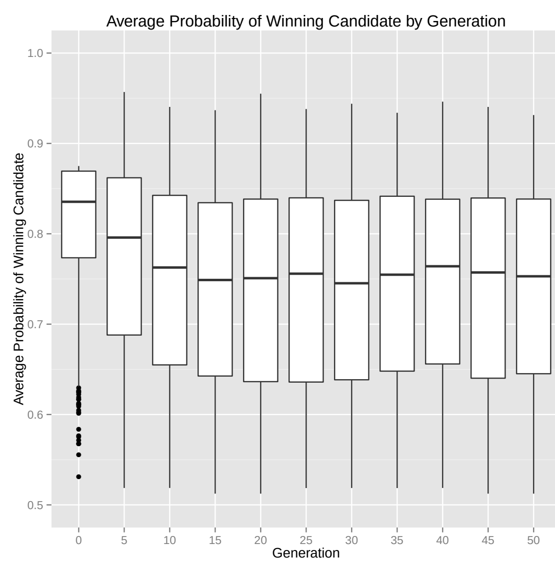

Figure 4: Average probability on the winning candidate at the end of each generation, for 100 learning steps per generation, with initial constraint weights of zero for new agents

deterministic grammar states, for both the Zero-Init and Rand-Init conditions. This can be seen in Figure 5 for the Rand-Init condition, and in Figure 6 for the Zero-Init condition. As the simulation progresses through the generations of agents, there is a visible trend towards increasing the probability of the winning output candidates.

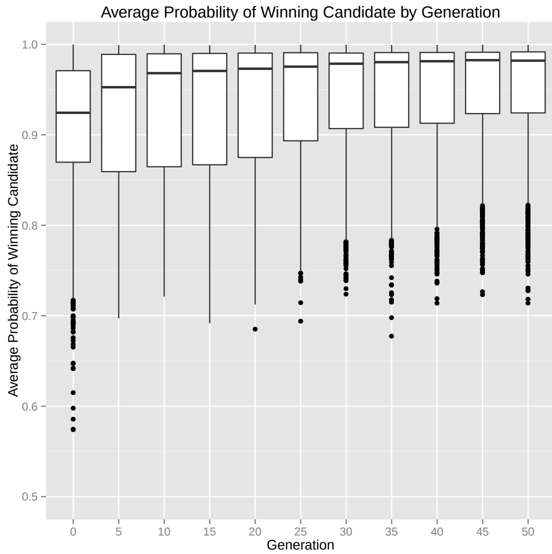

Figure 5: Average probability on the winning candidate at the end of each generation, for 1,000 learning steps per generation, with initial constraint weights of zero for new agents

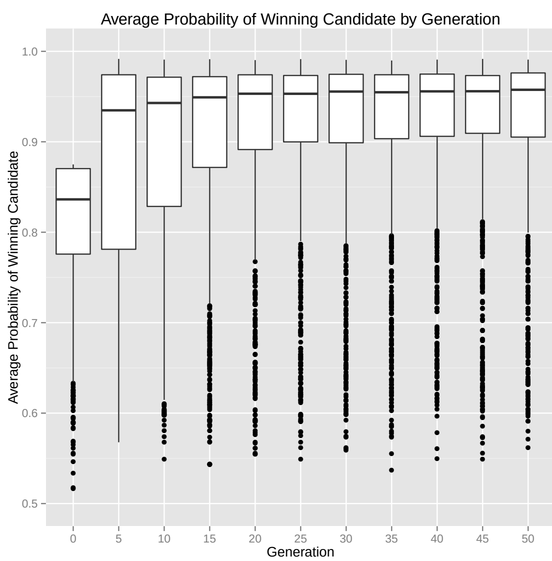

Figure 6: Average probability on the winning candidate at the end of each generation, for 1,000 learning steps per generation, with initial constraint weights of zero for new agents

5 Discussion

The results from the iterated learning model simulations are unlike the effects of the interactive learning model, which showed emergent biases towards deterministic grammar states and away from cumulative constraint interactions (gang effects) which were robust across parameter settings and agent initialization conditions. The iterated learning model results, on the other hand, demonstrated the same biases, but not in all cases: the bias away from gang effects emerged for both agent initialization conditions at higher learning step values (here, 1,000), while it only emerged at lower learning step values (here 100) when new agents were initialized with constraint weights of zero. The bias towards deterministic grammar states only emerged at higher learning step values, and not at all for lower learning step values.

The effect of both of these biases - away from gang effects and towards deterministic grammar states - means that overall these learning models produce ranked-constraint behavior from a weighted-constraint base grammar (MaxEnt). As a method of generating probabilistic typological predictions, the combination of weighted-constraint MaxEnt grammar and either of these agent-based learning models takes advantage of the greater representational power of MaxEnt while preserving some of the restrictiveness of ranked constraints through the influence of learnability differences between patterns in the typology.

References

Sharon Goldwater and Mark Johnson. 2003. Learning OT constraint rankings using a Maximum Entropy model. *Proceedings of the Workshop on Variation within Optimality Theory*, pages 113–122.

Coral Hughto and Joe Pater. 2017. Emergence of strict domination effects with weighted constraints. Talk given at the CLS Workshop on Dynamic Modeling in Phonetics and Phonology.

Simon Kirby. 2017. Culture and biology in the origins of linguistic structure. *Psychonomic Bulletin & Review*, 24(1):118–137.

Joe Pater. 2012. Emergent systemic simplicity (and complexity). *McGill Working Papers in Linguistics*, 22(1).

Juliet Stanton. 2016. Learnability shapes typology: The case of the midpoint pathology. *Language*, 92:753–791.

Robert Staubs. 2014. *Computational Modeling of Learning Biases in Stress Typology*. Ph.D. thesis.