Jeremy Needle

EECS 472 ABMs

U. Wilensky

9 June 2013

Language drift patterns from simple statistical learning in an agent-based model

Introduction

A fundamental question in linguistics is, 'by what mechanism do language shifts arise and spread?' We further seek to investigate the relative importance of the many hypothesized factors in this process. These factors include dozens of variables such as macro-scale social structures (e.g., race, class) and social network patterns, micro-scale social behaviors, cultural differences, mass media, etc. The investigation of this question is known as the Actuation Problem (Chen & Wang 1975), which can be expressed quite simply (Baker 2008): "The question raised by the actuation problem is simple: why does sound change occur?"

Language shifts are certainly know to occur; human languages do not remain static. These shifts occur time scales ranging from shorter than months, to longer than decades; and across every aspect of the language, including syntax, semantics, pragmatics, and phonology. Shifts also occur despite competing processes and principles (e.g., vowels show a tendency toward maximal dispersion within the vowel space; and communication obviously depends on a certain level of language stability) (Ohala 1989; Eckert 2000).

At a very basic level, human language can be understood to progress across generations: children learn from their elders, mature, and become elders themselves. Though this account leaves out vast areas of complexity, it may be a useful level at which to model this process. With such a simple model, observed historical trajectories of language may be approximated, which speaks to the validity of this basic generational assumption, and to the relative necessity of a more complex model.

Children are also known to use statistical learning (Saffran, Aslin, & Newport 1996). There are many possible mechanisms, ranging from the very simple (such as a naive Bayes classifier) to more complex constructs (e.g., exemplar theory). In addition, both perception and production are known to be noisy; therefore, a model that uses noisy statistical learning is appropriate and grounded in the literature.

Method

The basic generational mechanism models an ‘adult’ cohort speaking, and a ‘child’ cohort listening. Generations are iterated: at the end of each generational step (speaking/listening phase), the ‘adults’ die, and the ‘children’ select their language based on their listening; the children then become the new ‘adults’. The overall status of the language is monitored over many iterations. The models presented here build on this approach by probabilistically assigning adult agents a lifespan, which removes the mechanism of full-generation changes at each tick. Instead, some number of adults will die on a given tick, and only that number is replaced with newly ‘graduated’ children. The effect of varying the maximum lifespan is considered as a variable in this analysis. There are a number of other relevant variables manipulated in this model: Locality, Noise Level, and Learner Method. Locality refers to constraints placed on which adults children may listen to: instead of simply selecting from all adults, there are locality constraints that limit children to a random N subset of the adults, to the nearest-N adults, or to all adults within a certain N radius of the child. The secondary model instead places the agents on either a clustered or lattice network, within which they listen to their immediate (tie-distance = 1) neighbors.

Noise Level is simply how large the effect of noise on the learning process is, as a percentage of the total vowel space. Learner Method controls the statistical approach used by children in acquiring their language: mean takes the mean of the speaking group, median chooses the median values, and choose-one copies the value of one adult from the speaking group at random.

In these models, the 'language' is represented by two values which correspond to the first and second formants (F1, F2) of a single vowel; formants are the resonant frequencies of the human vocal tract, and provide a good description of vocal articulation during speech. Several additional acoustic factors are important, but these capture the majority of vowel information in a very concise way, appropriate for this model. The typical bounds of variation for these formants are from 500 to 3500 (F2) and from 200 to 1200 (F1), and these values are used in the model (the initial adults are seeded with these values according to a normal distribution).

Results and Validation

Validation

For these models, the first stage of validation means reliably producing vowel shifts across short and long time scales. In real life, we see vowel shifts across dialects, languages, and eras. Notable examples in English are the Great Vowel Shift (spanning over 300 years), the Northern Cities Shift (about 50 years), or the Southern Shift (about 30 years). For validation, the model must at least produce some shifts (defined as apparently sustained movement in population-language). The time scales of these shifts is also important: in real life, a shift may be on the order of generation(s), or much longer. For validation, the model should produce shifts across short and long time frames. Further validation will consider more precise aspects of the nature of the shifts: distance, speed, duration, trajectory patterns, etc.

Results

Shift patterns resulted across nearly all experiment treatments. Figures 1, 2, 3 show examples of what these shift patterns look like (Figure 1 shows a trajectory that mostly lacks any shifts). Note that the time frame is 3000 Ticks, which corresponds in the model to between 3000 and 600 generations; even quite short periods of excursion could be interpreted by humans living through them as ‘shifts’.

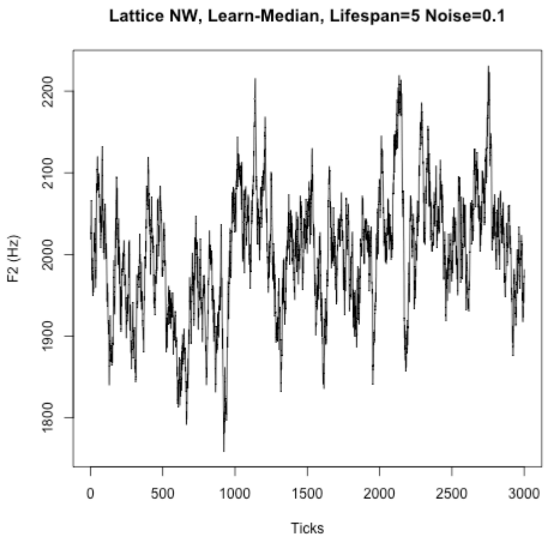

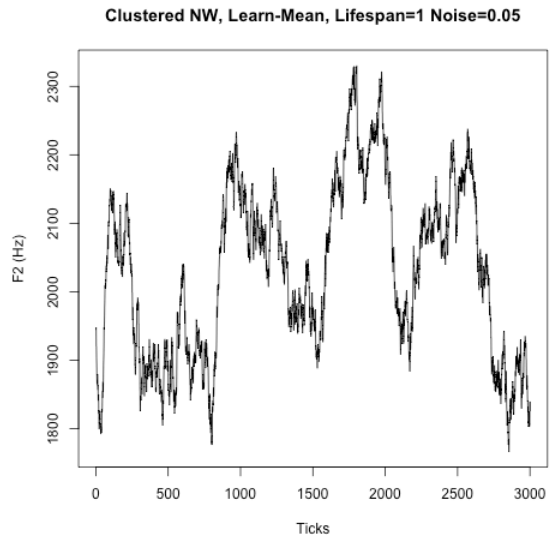

Figure 1, Figure 2: Formant trajectories showing no-shift, long-shift

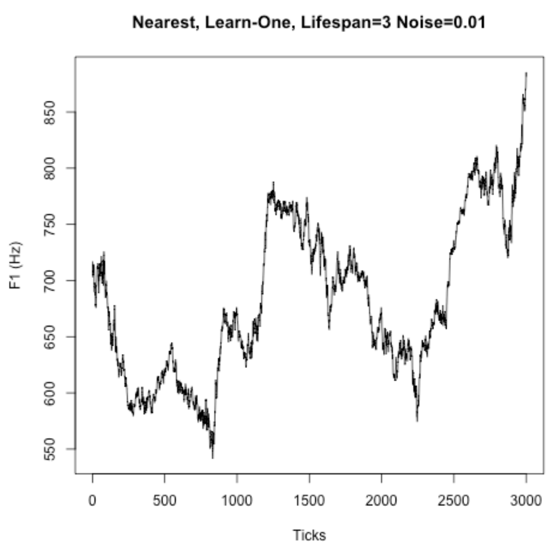

Figure 3: Formant trajectory showing shift

Figure 4 and 5 show that shifts are present across all tested Locality categories, Noise levels, and Lifespans; for certain combinations (e.g., ClusterNW with 10% Noise), shifts are less common. In general, the ClusterNW was most likely to inhibit shift presence, followed somewhat by the LatticeNW locality; other localities did not inhibit shifts. Noise above 10% also inhibited shift presence in some cases, because the noise was so strong that it overwhelmed the (weak) shift itself. Though it was hypothesized that longer lifespans might exert a stabilizing effect on the language population, this effect did not materialize: Lifespan did not appear to play a significant role in shift presence. The presence of ‘long’ shifts was also evaluated; a long shift is one which was stable in its new position for more than 10% of the time span (i.e., more than 300 Ticks). These effects are shown below in Figures 6 and 7.

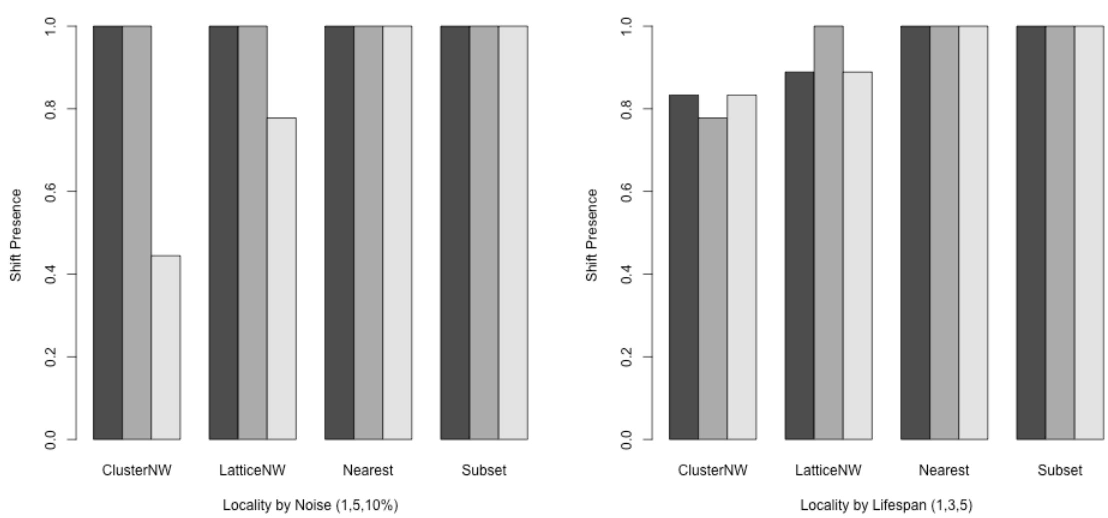

Figure 4, Figure 5: Shift Presence for Locality by Noise, Lifespan

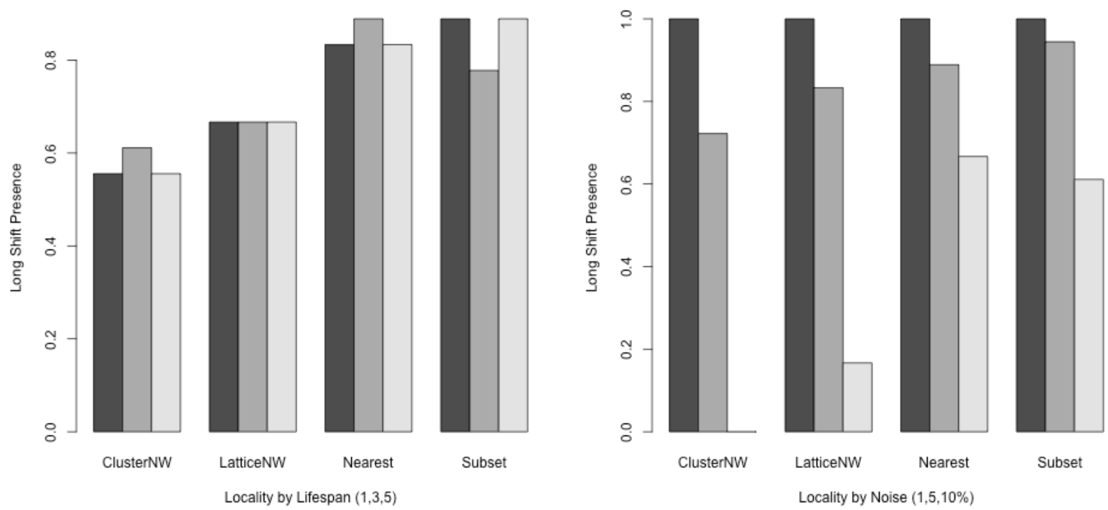

Figure 6, Figure 7: Long Shift Presence for Locality by Noise, Lifespan

The effect of the different simple statistical learning strategies was also tested. As shown in Figures 8 though 11, there does not seem to be significant differences in shift or long shift presence based on learner type.

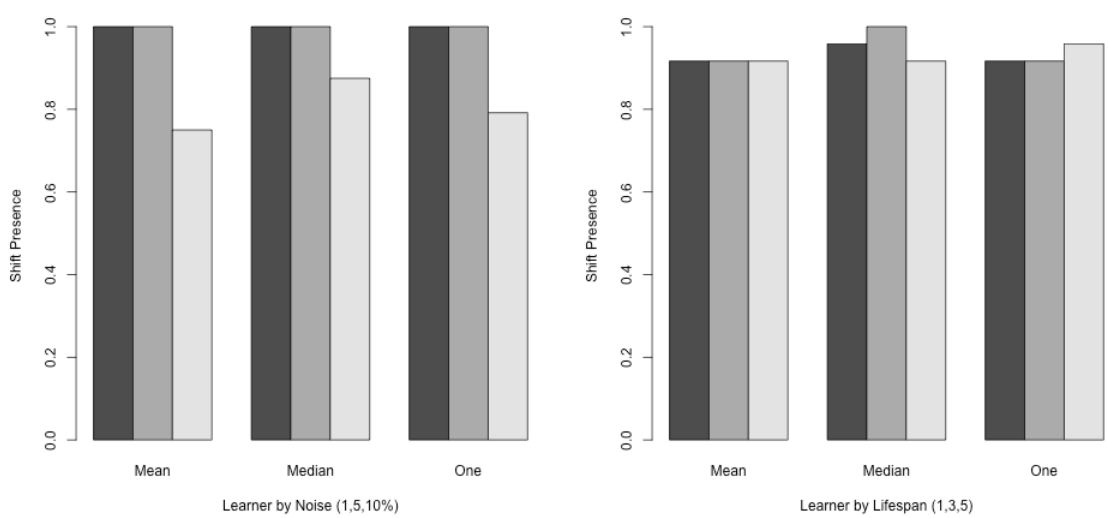

Figure 8, Figure 9: Shift Presence for Learner by Noise, Lifespan

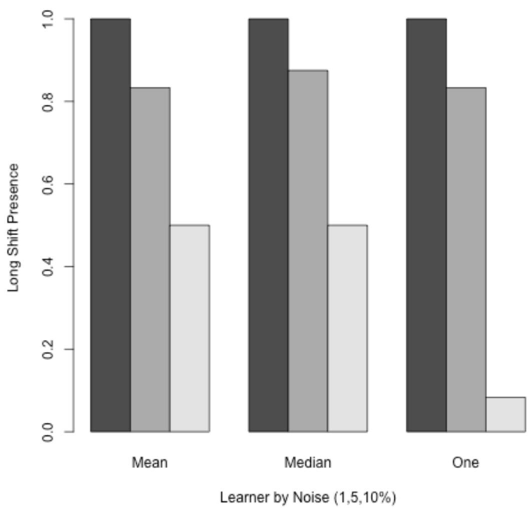

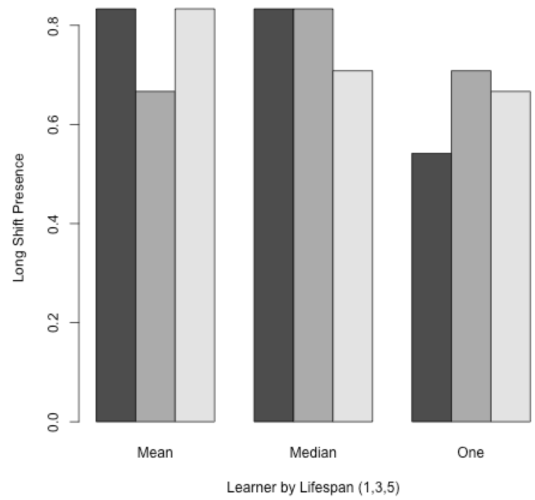

Figure 10, Figure 11: Long Shift Presence for Learner by Noise, Lifespan

Discussion and Limitations

The results show that the initial stage of validation is successful, with the simple models producing sustained shifts at appropriate time-scales across nearly all combinations of parameter settings. This provides evidence that the answer to ‘why do language shifts occur?’ is not as complex as previously hypothesized; instead of involving the interplay of many individual and social factors, a simple noisy statistical learning model is sufficient for shifts to arise. Note also that a ‘shift’ may simply be a sort of smoothed random walk, when the observers are only taking into account generation-level time scales: the walk is simply too slow for the speakers to notice its nature.

Limitations

It is true, however, that these models only captures some of the nature of apparent shifts. Other factors are known to have effects on language development: culture, social network structure, migration, language contact, social structure and social class, etc., all exert their own influences on the trajectory of language.

In addition, these models only consider the change in a single gradient phonetic feature (a idealized ‘vowel’), which means they do not necessarily model the language change trajectories taking place in the lexicon, syntax, pragmatics, semantics, and the rest of the phonological system. These model also avoid some effects that have direct analogues for this speaker/listener model, such as the presence of multiple vowels, effects of population size, heredity/family structure, or gender roles.

Extension: Network Model

The secondary model attempts to reproduce the results of the main model while exploring network structure effects on those results. The agents are remapped onto either a random locally clustered network (with density k), or onto a 2D lattice (with link probability p; no isolates allowed). The results of this secondary model are discussed above (Locality = ClusteredNW, LatticeNW). However, the is ample opportunity for further work with this model. The first followup will be the incorporation of the NetLogo NW extension, and the corresponding implementation of further network structures, such as Preferential Attachment. Another option is to import a realistic social network structure directly from observational data.

Future Work

The next step is to refine the models used in this experiment, and to accordingly refine the statistical methods used for validation. With more precise time-series analysis methods, it should be possible to quantify the degree to which the experimental vowel trajectories differ from each other, yielding a better understanding of the influence of each variable tested.

References

Baker, A. (2008). Addressing the actuation problem with quantitative models of sound change. University of Pennsylvania Working Papers in Linguistics, 14(1), 3.

Chen, M. Y., & SY, W. (1975). Sound change: actuation and implementation. Language, 255-281.

Eckert, P. (2000). Language variation as social practice: The linguistic construction of identity in Belten High. Blackwell.

Ohala, J. J. (1989). Sound change is drawn from a pool of synchronic variation. Language change: Contributions to the study of its causes, 173-198.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. Science, 274(5294), 1926-1928.