Emergent Behavior in Phonological Pattern Change

Mark Dras[1] and **K. David Harrison**[2]

[1] Department of Computing, Macquarie University, North Ryde 2109 Australia (`madras@ics.mq.edu.au`)
[2] Department of Linguistics, Swarthmore College, Swarthmore PA 19081 USA (`dharris2@swarthmore.edu`)

Abstract

Language change has recently come to be seen as a complex dynamical system, along the lines of evolutionary biology and economics, as opposed to previous conceptions as a linear or cyclical system. We model the change of a particular phenomenon, vowel harmony, and look at the conditions under which the trajectory of change matches theoretical and empirical predictions. Our experimental work shows that there are certain conditions under which the desired trajectories do not occur, implying that absence of these conditions is necessary for accurate modeling of language change.

Introduction

Language change has at various times been seen as linear—that is, languages are progressing or decaying monotonically—or cyclical—that is, languages pass through a life cycle of birth, maturity, death and rebirth. However, modeling language change in a formal way has led to a recognition that it is a complex dynamical system (Lass 1997): the interaction of individual speakers leads to emergent, global population characteristics of a language that are neither linear nor cyclical.

In all dynamical systems there is an observable pattern of change, and the object of a model is to describe the process such that the model output closely approximates the empirical data; e.g. in evolutionary biology, there is gene frequency change, the basic process of evolution. In language change, there is the replacement of one linguistic phenomenon by another, such as when historically English diverged from Germanic and lost the requirement that the verb come second in the sentence. Analogy with evolutionary biology (Bailey 1973) and empirical work (Chen & Wang 1975) have supported an S-shaped functional dependence on time as the fundamental trajectory of language change.

The focus of our research is vowel harmony in Turkic languages. These systems exhibit a great deal of change—instability, even—over the more than one millennium during which these languages have been recorded in writing. Most Turkic languages have two autonomous harmony systems, one based on tongue backness and one on lip rounding. The systems range from

	front	back
high	i y	ɯ u
non-high	e ø	a o

Table 1: Vowel inventory (IPA)

robust, nearly exceptionless harmony to highly variable or restricted harmony to no harmony at all. They thus provide rich material for constructing typologies of harmony and models of language change.

Changes at the level of the individual aggregate to give language-wide evolution in backness harmony systems. Agent-based simulation, where a community of computational agents interacts over time, is an appropriate way of modeling language-related complex systems, as language changes take place in a social context. In this paper we present some results on what sorts of simulation models can produce observed trajectories of language change, and which cannot.

Vowel harmony

In most Turkic languages, backness harmony is apparent both as an ambient pattern of vowel co-occurrence within word roots, and as a productive pattern of vowel alternations (e.g., in suffixes). The condition imposed by harmony is that only vowels from the same class can co-occur in a particular context.

A typical Turkic vowel inventory includes four front and four back vowels, divisible into harmonic classes (Table 1). Under backness harmony, then, only back vowels can co-occur with other back vowels, and only front vowels with front vowels. As an example of the first manifestation of backness harmony, word roots can only contain all back or all front vowels; this is particularly apparent when words are adopted into a harmonic language from a nonharmonic one (loanwords), where these are mutated so that the localized word is harmonic. For example, the disharmonic word *araki* ('alcohol') in Mongolian, a non-Turkic language, was adopted into Tuvan, a Turkic language, as *araɣa*. Similarly, the French 'chauffeur' (phonetically ʃofør) has been adopted into some dialects of Turkish as ʃøføør. As a manifestation of the second

type of vowel harmony, regarding suffixes, the suffix must match the vowel in the root it attaches to. For example, in Turkish *at* ('name') takes *-lar* to become the plural *atlar*, but *ev* ('house') takes *-ler* to become *evler*.

There are a number of hypotheses about how backness harmony may have emerged: that harmony arises from co-articulation, where the shape of the mouth from previous vowels predisposes the speaker towards uttering another with the same mouth shape; the structuralist notion that symmetry in vowel inventories provides an impetus to harmony; etc. (see Harrison and Dras (2001) for further references).

Historically, Turkic vowel harmony systems are constantly in flux. Old Turkic as attested in 8th-11th century runic inscriptions from Siberia had an eight vowel system and fully regular backness harmony. Modern Turkic languages have from 5 to 10 vowels, and range from almost fully harmonic (Tuvan) to not harmonic at all (Uzbek). The Turkic family thus provides over one millennium of documented stages and scenarios in the evolution of harmony. These serve as data points showing harmony evolution along a definable trajectory.

But not all points on this trajectory are discernable in the historical record. There are stages we know must have taken place that were not recorded. For example, 8th century Old Turkic shows pervasive, almost exceptionless vowel harmony for backness. Prior stages in the emergence of this system were not documented, and the gap limits our empirical knowledge of how such systems originated. By contrast, the evolution of harmony systems in daughter languages of Old Turkic is quite well documented across a period of over 1,000 years, allowing us to precisely quantify stable or declining levels of harmony over time, up to the present day.

Related Work

Specifically on phonology, de Boer (2000) looks at how vowel systems can arise from nothing and how the vowels organise themselves in the vowel space of the population of speakers. He notes, however, that his aim is not to model historical evolution of vowel systems, because of its greater complexity; one such example of this type of complexity is the situation when there is structure (like harmony relations) within the vowel space.

We are interested in modeling the historical evolution of vowel systems; the work in this paper is one step towards understanding the sources of this extra complexity. Specifically, we look at the historical trajectory of vowel harmony evolution. Other work on simulation modeling looks almost exclusively at the binary question: Does the phenomenon emerge (or decay) at all? For example, Steels (1997) asks whether a shared vocabulary can emerge from agent interaction; Kirby & Hurford (2002) whether syntax can emerge in a simulation from initial randomness; Zuraw (2001) the more specific

question of whether nasal coalescence in simulated Tagalog will replicate real life. In these situations there are effectively only two outputs to compare with observations: the start point (phenomenon not present) and the end point (phenomenon either present or not). Having such a small set for comparison does not allow one to determine with much confidence what the factors causing the phenomenon are: many possible factors could cause the same behavior. Modeling change as a trajectory constrains the simulation to a much greater degree, allowing us to rule out many possibilities. While it is of course possible to evaluate binary models by systematically modifying parameter settings, there is no corresponding empirical data to match it. That is, there is typically not a range of sets of linguistic data under minimally different conditions apart from changes over time. Wonnacott (2000) has also begun to look at how S-shaped curves as a trajectory of change arise, although in that work the factors that distinguish between full S-curves (slow–fast–slow change) and partial S-curves (an exponential fast–slow change) are not examined. In this paper we aim to tease out some such factors.

In building our simulation we adopted the full S-curve as a trajectory of evolutionary change. Several independent lines of research suggest that language change often proceeds along this trajectory. The rising curve shows the advancement of a new form at the expense of an old one. On the curve, change begins slowly, accelerates, then slows again, over a period of many generations. It was originally proposed for language change in Bailey (1973), as part of a "wave" model of linguistic change, with support coming from parallel behavior in population biology in the replacement of genetic alleles. As empirical support, Chen and Wang (1975) look at three case studies of historical data that demonstrate the S-shaped behavior of language change: the Chaozhou dialect of Chinese, where words have been shifting from one tone class to another, with the slow-fast-slow pattern in evidence; English diatones; and the Swedish optional final *-d*. The earliest work on quantitatively modeling language change (Kroch 1989), on the transition of Old English and Old French away from verb-second syntax, thus adopts the S-shaped curve, as has subsequent work, but in the model imposes it on the data. Kroch proposes:

> ... given the mathematical simplicity and widespread use of the logistic [a particular equation giving an S-curve], its use in the study of language change seems justified, even though, unlike in the population genetic case, no mechanism of change has yet been proposed from which the logistic form can be deduced. (Kroch 1989: 204)

Later work, including our own, is interested in how the S-shape observed in data can emerge from simple parameter interaction. The first of these are macro models that model the behavior of the whole speech community through mathematical recurrence relations

(Niyogi & Berwick 1997). So, for example, the proportion of the population that is using the new variant at some time n, p_n, is a function of p_{n-1}, of the form $p_n = Ap_{n-1}^2 + Bp_{n-1} + C$, where A, B and C are coefficients determined by a model of language acquisition. However, these models have fundamental problems because they treat populations in the aggregate, and moreover non-stochastically (Briscoe 2000); Niyogi and Berwick's model in particular produces implausible equilibria under certain conditions. Incorporating stochastic behavior, and subdividing the aggregated population, leads logically to a computational agent-based simulation, as the mathematics otherwise becomes intractable.

In the case of Turkic harmony emergence, we are assuming an S-curve trajectory in the absence of historical data points. In the case of harmony breakdown, we also adopt an S-shaped curve, but we are guided here by a number of historical data points along the trajectory. We are not claiming that the S-curve is necessarily the right curve, merely that it is a plausible one for this type of change. For now, the S-curve is more or less in accord with the evolution we have been able to map out for Turkic harmony. It could be suggested that an S-curve will be the result of almost any model that is proposed. However, we demonstrate that that is not the case, and use it to rule out one class of possible models.

Basic Model

An important principle in building the simulation—in choosing which factors to include in the model, and in choosing how to realize them—is, following Occams Razor, to start with as simple a model as possible. If this fails to model the data accurately, the model is made incrementally more complex until (hopefully) it properly fits the data. If we start with a complex model, it isn't possible to tell which factors are crucial for the outcome.

Globally in the simulation there is the language, consisting of 1000 disyllabic words based on a real Turkish lexicon. The proportion of harmonic words is a parameter, by default 50%—for disyllabic words, if front and back vowels are (overall) equally likely, this is the mean level of harmony that would occur just by chance. The words are represented by strings in the implementation, with vowel phonemes separate symbols with associated features (backness, height, roundedness).

An individual agent has a lexicon that is a subset of this language, and is capable of speaking, listening and reproducing. It is only aware of its own neighborhood, defined as the four adjacent spaces. Agents do not move. Throughout the simulation agents are distributed with medium density, so that on average an agent will have two neighbors.

An agent lives for between 20 and 50 turns[1]. It reproduces only once, so the population is stable. It begins

[1]All random values are from a uniform distribution

with a starting lexicon of 6 words. For the initial population, this is taken from the global language; for others, it is taken randomly from the parent's lexicon. The lexicon also grows through conversation with a neighboring agent; a parameter controls the likelihood of a conversation taking place in a given turn.

A conversation consists of a single random word, one agent speaking and one listening. In the conversation, there is a uniform probability for all agents that the word will mutate either towards or away from harmony through the various factors above (co-articulation, misperception, hypercorrection); there are different probabilities for harmonizing and for disharmonizing. In a mutation, one vowel in the word is altered by changing the polarity of the backness feature.

This is so far very simple. We could, in principle, use a more complex mutation, for example, one that incorporates Bayesian probabilistic reasoning, as modeled in Zuraw (2001). However, it is not clear that this level of complexity is necessary—in fact, alternative models for Zuraw's results could be proposed using simple uniform probabilities that explain the data equally well. We show later with our own results that we can model the trajectory successfully without Bayesian probabilities.

We then plot a curve representing the proportion of harmony in the language. This is the average across all agents of the harmony of each agent's lexicon (that is, the proportion of words in the lexicon that are harmonic). This is not the only possible measure of language harmony, but it does reflect both the change in number of people adopting more extensive harmony, and the number of words that have become harmonized. Our first result is as in Figure 1. In this curve, the probability of harmonizing is 0.3 and of disharmonizing is 0 (in keeping with our native speaker informant's intuition that words are never deharmonized). The curve rises, and eventually plateaus at 1 (i.e. universal harmony), but too steeply: there is no period of slow change at the start. It is as if the agents were obeying a directive to "go forth and harmonize", rather than harmony evolving organically from social interactions: the result looks like a typical curve seen in machine learning. The same behavior occurs under all parameter settings tried where the probability of disharmonizing is 0; when this latter probability is non-zero, there is only random oscillation. Thus it is not the case that any type of model of interacting agents will produce a full S-shaped curve of some arbitrary phenomenon.

Variant Model

Systematically varying parameters never generated a real S-curve under the basic model. We therefore moved to a slightly more sophisticated definition for the probability of word change. Here we recognize that not all people in the real world will be equally likely to mod-
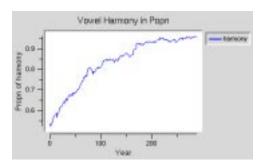
Figure 1: The basic model

Figure 2: The extended model

ify a word. For example, in adopting 'chauffeur' from French, a speaker from Istanbul is more likely to keep close to the original vowel sounds ($\int of \emptyset r$), whereas a villager from eastern Turkey whose lexicon contains many fewer disharmonic (foreign) words may harmonize it, to $\int \emptyset f \emptyset r$ (in fact, both variants are attested in colloquial Turkish). That is, the likelihood of harmonizing a word is correlated with the strength of the pattern of harmony. We model this by making the probability of word change linearly proportional to the proportion of harmony in the lexicon. To do this, we replace the previous parameter representing the given probability of word mutation with a maximum probability of word mutation, and an agent's individual probability will be some fraction of this.

In addition, recognizing that a "pattern" that covers, say, 2% of the lexicon is unlikely to be pervasive enough to be really considered a pattern for an individual, we can set a threshold T below which a pattern is not recognized as significant. So for an agent x:

$$
\begin{aligned}
h(x) &= \text{proportion of } x\text{'s lexicon that is harmonic} \\
ph(x) &= \text{MPH} * h(x)
\end{aligned}
$$

where MPH is a parameter representing the probability that an agent with a fully harmonic lexicon would mutate a word. Then the probability of word mutation given the harmony of the lexicon, P, is

$$
P = \begin{cases} \text{ph(x)} & \text{if } ph(x) > \text{T} \\ 0 & \text{otherwise} \end{cases}
$$

In a sense, agents behave as if they assume that their neighbor's grammar resembles their own, even if they don't know exactly what it is.

In contrast to the basic model, this produces the first full S-curve, as in Figure 2. For this particular curve, MPH is 0.9 and T is 0.65. Smoother S-curves can be generated by lowering the threshold, with the middle section tending more towards linear the lower the threshold; at higher thresholds the curve is less clear, and could possibly be random drift followed by exponential change after passing a bifurcation threshold.

Conclusion

Our conclusion is that we have to some extent narrowed down how the S-shape in language change arises. Un-
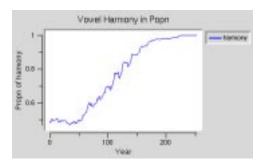
like earlier models, we are able to specify properties of individuals that lead, without explicit programming, to the emergent population behavior at least as described by its trajectory. In particular, we found that the most simplistic combinations of factors do not lead to the desired trajectory. However, if an agent does take into account its existing pattern of harmony in evaluating new words, effectively believing that its neighbors are similar to itself by projecting its own pattern onto other agents, it is possible to generate the S-curve of vowel harmony emergence. This incremental modification of the model, to one that produces the upward S-curve, gives us confidence that the model is a starting point for answering linguistic questions. Future work is to look at factors that could lead to a downward S-curve, such as homophony, and test against actual historical data.

References

Bailey, C.-J. 1973. *Variation and linguistic theory*. Washington DC: Center for Applied Linguistics.

Briscoe, T. 2000. Macro and micro models of linguistic evolution. In *Proceedings of the 3rd International Conference on Language and Evolution*.

Chen, M., and Wang, W. 1975. Sound change: actuation and implementation. *Language* 51(2):255–281.

de Boer, B. 2000. Self organization in vowel systems. *Journal of Phonetics* 28(4):441–465.

Harrison, D., and Dras, M. 2001. Agent-based modeling and microparametric variation: Modeling the evolution of vowel harmony. In *Proc. of NELS32*.

Kirby, S., and Hurford, J. 2002. The emergence of linguistic structure: An overview of the iterated learning model. In Cangelosi, A., and Parisi, D., eds., *Simulating the Evolution of Language*. Springer Verlag.

Kroch, A. 1989. Reflexes of grammar in patterns of language change. *Language Variation and Change* 1:199–244.

Lass, R. 1997. *Historical Linguistics and Language Change*. Cambridge University Press.

Niyogi, P., and Berwick, R. 1997. Evolutionary consequences of language learning. *Linguistics and Philosophy* 20:697–719.

Steels, L. 1997. The synthetic modeling of language origins. *Evolution of Communication* 1:1–34.

Wonnacott, E. 2000. Investigating the trajectory of language change. Master's thesis, University of Edinburgh.

Zuraw, K. 2001. The extension of semiproductive morphophonemic processes to novel forms: An optimality-theoretic model and simulation. In *Proc. of WCCFL XX*.