

## The new big picture for linguistics: Complex systems

William A. Kretzschmar, Jr.

*kretzsch@uga.edu*

*University of Georgia*

### Abstract

In the history of linguistics there have been crucial moments when those of us interested in language have essentially changed the way we study our subject. We stand now at such a moment. In this presentation I will review the history of linguistics in order to highlight some past important changes in the field, and then turn to where we stand now. Some things that we thought we knew have turned out not to be true, like the systematic, logical nature of languages. Other things that we had not suspected, like a universal underlying emergent pattern for all the features of a language, are now evident. This emergent pattern is fractal, that is, we can observe the same distributional pattern in frequency profiles for linguistic variants at every level of scale in our analysis. We also have hints that time, as the persistence of a preference for particular variants of features, is a much more important part of our language than we had previously believed. We need to explore the new realities of language as we now understand them, chief among them the idea that patterned variation, not logical system, is the central factor in human speech. In order to account for what we now understand, we need to get used to new methods of study and presentation, and place new emphasis on different communities and groups of speakers. Because the underlying pattern of language is fractal, we need to examine the habits of every group of speakers at every location for themselves, as opposed to our previous emphasis on overall grammars. We need to make our studies much more local, as opposed to global. We do still want to make grammars and to understand language in global terms, but such generalizations need to follow from what we can now see as the pattern of language as it is actually used.

**Keywords:** complex systems; A-curve; emergence; alternative facts; Big Data.

A development in 2017 that has had a striking affect in America is the new idea of “alternative facts” or “fake news.” These two labels actually go together. We have all experienced the information explosion on the Internet, which is mostly a good thing but which also raises the problem of authority.

Which information should we believe? It used to be that the publication process, the fact that somebody had to pay money to publish information and that we ourselves paid money to acquire it, offered some suggestion of authority. Publishers and editors have insisted on checking facts. In academic writing, the process has involved peer review. Now, however, anybody can put anything they want on the Internet, and we can see what they put up, without any very obvious cost to the poster or the reader. Without peer review or editorial boards, Internet information is not authoritative. Again, this is mostly harmless, but it does put the burden of deciding what to believe on the reader, who can no longer rely on the publication process as a guarantee. Thus “alternative facts” and “fake news.” Readers can decide that one set of facts they see on the Internet is what they want to believe, and another set of facts is not. Readers can decide that one news report is what they want to believe, and another report is “fake news.” The “facts” and the “news” are therefore no longer grounded by what actually happened.

They are just information, and consumers get to decide what information they like the best. Some readers get involved with fact checking, to try to determine whether the information on the Internet is evidence-based, but so far the fact checking has not carried the day because it does not have any special credibility in the new information space. Fact checking only convinces people who want to believe it.

What I would like to argue is that linguistics has been too much like “alternative facts” and “fake news”. The “facts” and the “news” about languages as described by linguists have not been grounded by what all the speakers of a language actually say and write. That said, I want to emphasize that I am not accusing linguists of political adventurism! Linguists have had reasons for what they have done. Nonetheless, what we have considered to be facts about language, and news about linguistic change, have not had the same kind of empirical scientific adequacy that is required in the natural and physical sciences. This essay will describe how this came to be, and then will discuss how we can establish empirical scientific adequacy for linguistics. In a word (actually two words) complex systems is the scientific model that fits what people actually say and write. This is the new Big Picture for linguistics today: we need to use new tools and methods to make a new linguistic science for the twenty-first century.

First, a look back into the history of how we have studied language, and how linguistics as a science came to be. Before the nineteenth century, the study of language generally meant the study of Latin, particularly its gram-

mar. Newton's *Principia*, for instance, was written in Latin. Early dictionaries and grammars eventually appeared for languages other than Latin, but most people thought of language not as something systematic but as something that concealed a spiritual power. The Gospel of John, in its seventeenth-century King James translation, said that "In the beginning was the Word." Magic, another mystical pursuit, involved the use of special words in spells, as we still see today in the world of Harry Potter. However, in the nineteenth century, new linguists, first comparative and historical analysts and later the NeoGrammarians, brought new systematic study to language. They used evidence from dictionaries to understand how languages might be related and how they might have changed. The NeoGrammarians wanted change to be mechanical and exceptionless, in line with the mechanical sciences of their time. However, when Georg Wenker tried to bring this mechanical theory to German dialects, he found from the beginning of his survey research that what people said was more irregular than the evidence of dictionaries might suggest. As he said in 1885, "I lived in the fair and calming conviction that these [linguistic] features must completely or nearly completely go together. That assumption turned out soon enough to be utterly mistaken: the boundaries of the contemplated features stubbornly took their own way and often crossed each other" (see Kretzschmar 2009, Ch. 2). So, by the end of the nineteenth century, the study of language had become systematic, but language itself as Wenker surveyed it was not a mechanical system.

At the beginning of the twentieth century, Saussure set out to define linguistics as its own special discipline. The only way that Saussure could do that was to assert that (Kretzschmar 2009, Ch. 2):

The linguist must take the study of linguistic structure as his primary concern, and relate all other manifestations of language to it. Indeed, among so many dualities, linguistic structure seems to be the one thing that is independently definable and provides something our minds can satisfactorily grasp.

In Saussure's view, linguistic structure arose from seeing language as a "collective phenomenon" while, on the other hand, "there is nothing collective about speech. Its manifestations are individual and ephemeral. It is no more than an aggregate of particular cases. ... Language in its totality is unknowable, for it lacks homogeneity" (Kretzschmar 2009, Ch. 2). The evidence of speech, people's talk as Wenker had surveyed it, did not have the homogeneity

ty, the unity, that linguistic structure had. The unity of linguistic structure is imposed upon it by the analyst: it is *definable*. The linguist makes a generalization from particular cases of speech. Still, Saussure says that, while linguistic structure is necessary, speech is always prior (Kretzschmar 2009, Ch. 2):

speech also is necessary in order that a language may be established. Historically, speech always takes precedence. How would we ever come to associate an idea with a verbal sound pattern, if we did not first of all grasp this association in an act of speech? Furthermore, it is by listening to others that we learn our native language. A language accumulates in our brain only as the result of countless experiences. Finally, it is speech which causes a language to evolve. The impressions received from listening to others modify our own linguistic habits.

In Saussure's view, linguistic structure depends upon speech both for the linguist and for individual speakers. While speech may be individual and ephemeral, it is the only way to get to the definable generalization of linguistic structure.

This idea was in the air at the time from Durkheim's invention of sociology, as J. R. Firth has pointed out (Kretzschmar 2009, Ch. 5):

De Saussure's general linguistics is closely linked with the sociology of Durkheim. His theoretical approach may fairly be described as Durkheimian structuralism. ... De Saussure, thinking in Durkheimian terms, regarded social facts as *sui generis* and external to and on a different plane from individual phenomena. ... The group constrains the individual, and the group culture determines a great deal of his humanity.

The idea of a "social fact" thus takes a generalization defined by the linguist, linguistic structure, and turns it around so that the social fact constrains the behavior of individual speakers. Indeed, as Firth argued,

true Saussureans, like true Durkheimians, regard the structures formulated by linguistics or sociology as *in rebus*. The structure is existent and is treated as a thing. As Durkheim said, such social facts must be regarded "comme des choses."

So, linguistic structure can lose its status as a definable generalization made by linguists, and become instead something "treated as a thing" with its own

reality. While Noam Chomsky cannot exactly be called a “true Saussurean,” we can understand in Saussurean terms his treatment of language as a biological phenomenon, somehow encoded in our genes, and his rejection of behavioral evidence like the collection of corpus data. For Chomsky, language is a biological, genetic thing, *in rebus* just as Firth described Saussurean ideas following Durkheim. Perhaps we should not be surprised that we have not found very good evidence that the social fact of language, linguistic structure defined by linguists from their experience with speakers, has become a biological, genetic thing. Numerous arguments, for example by Tomasello (2003, as cited in Kretzschmar 2015, Ch. 4), discuss that issue in detail. At this point we can return to Saussure, and in particular to his assertion that “language in its totality is unknowable, for it lacks homogeneity.” While that was true at the turn of the last century, it is no longer true today. When Saussure was writing it was not possible to preserve language or to store and manipulate it. He did not have effective methods of survey research. Today, however, we do have computers to store and manipulate evidence of speech. We do have survey research methods that allow us to estimate the “totality” of language in a community. In short, we can now use sampling and Big Data. Moreover, when we do these things, we see that language does indeed have a sort of “homogeneity.” It turns out, when we collect evidence of any linguistic feature, whether lexical, phonetic, or grammatical, that we can observe the same quantitative frequency profile underlying all the different ways to say the same thing. This frequency profile occurs at every level of scale in our analysis, from an overall sample for a population down to every subsection, like just the women or just the men, or just the members of different occupations. Human speech is thus like many other aspects of the natural world, as discussed persuasively in Mandelbrot (1977, 1982): it is fractal.

The fractal structure of speech has two properties, nonlinear frequency distributions and self-similarity at every level of scale. We always see that the variants for some feature, like the variant realizations of the [æ] vowel (Figure 1, from my Eastern States survey data), will occur in a nonlinear pattern where a few types of realizations are very common, some types are moderately common, and most types are rare.

Let us call these distributions A-curves, short for asymptotic hyperbolic curves. They are selfsimilar at every scale (also known as “scale-free”), in that the same A-curve pattern will appear in every subset of the data, whether for the overall set, or for any regional subset, or for any social subset, or for

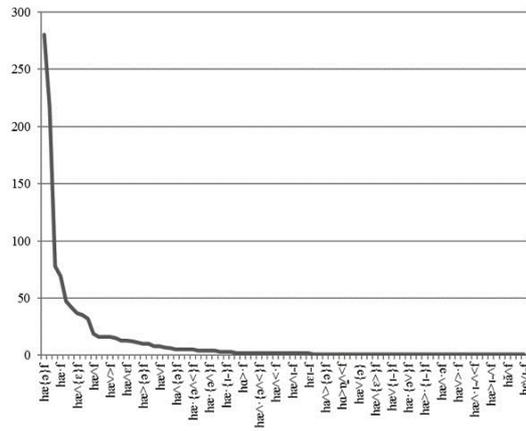


Figure 1. Variants of the [æ] vowel in *half*, from the Linguistic Atlas of the Middle and South Atlantic States ([www.lap.uga.edu](http://www.lap.uga.edu)).

any individual. The only way to observe such a “homogenous” pattern at every level of scale is to use a very large data set, Big Data. Of course this nonlinear distribution has been known for a long time. George Zipf popularized it in the first half of the twentieth century, and Zipf’s Law has been recognized as an oddity in linguistics since that time (for discussion of the literature, see Kretzschmar 2015, Ch. 4). What is different now is our knowledge from Big Data that the A-curve is not an oddity but a fundamental property of speech.

The underlying A-curve allows human speech to be both understandable and flexible. As Figure 2 shows for all the different terms for what to call a *thunderstorm*, the terms that are top ranked in the frequency profile provide the most understandable variants. This small number of variants accounts for 80% or 90% of what people say, here what people call thunderstorms. The rest of the variants are hardly ever used, and they provide flexibility in speech. Even though these variants in the long tail of the curve are rare, we can most often understand them in context, and they allow every speaker or group of speakers to mark themselves as different in their speech, to be themselves in their speech. Thus we can recognize individual speakers for themselves, and we can often recognize speakers as coming from some particular place or particular group.



read anything. Our minds are not like computers that rely on exact matches of text strings in searching and sorting data as if in a formal dictionary. We know, and computers do not know except for the special case of spell checkers, what word was meant when it appears with an odd spelling. We have some latitude in how we put stimuli together to make language, whether when we see letter shapes and word spellings, or when we hear pronunciations and understand words from those. As users of language, we are always reaching for the best fit of what we hear or see. Sometimes we do fail to put things together in the same way that our conversational partner has meant them, and these situations are often memorable for us, often quite funny because of the unanticipated broken expectations. But most of the time we do not fail to put things together. We use the underlying homogenous A-curve for every feature of language – pronunciation, words, grammar, discourse – in order to understand each other and to be ourselves. We could not communicate without it.

The A-curve pattern is the result of the complex system of human speech. Complex systems are an alternative model in science to the more common, Newtonian, cause-and-effect model. They are relevant wherever large numbers of elements interact with each other and patterns emerge from the interaction. There are no simple causes, but rather an accumulation of influences that, all together, allow an effect to emerge. Complex systems are thus relevant in quantum mechanics, as quanta interact and effects emerge; in our immune systems, as many cells interact and resistance to diseases emerges; and in economics, where many buyers and sellers interact in trades from which markets emerge. Human speech is thus not one off, but an example of a kind of interactive behavior that occurs in many domains in nature and society.

For any complex system to operate, we must have the random interaction of a large number of components. This interaction constitutes activity, which is essential to the operation of the system, for a complex system dies if the activity stops. The usual practice in linguistics has been to compose a static, hierarchical grammar, which is possible only as a snapshot of the complex systems process. Thus there is nothing wrong with grammars, but they do not represent the activity of the complex system which is one of its essential properties. As the interactions continue the components affect each other in what might be called an exchange of information, which leads to reinforcement of behaviors among the components, and thus to the emergence of stable patterns in the complex system, all without any central control. All of

these things happen in relation to contingency, to local circumstances, so that complex systems are able to react to circumstances, to change as they need to over time. The A-curve pattern in speech thus emerges from massive numbers of interactions between speakers. Just as in economic markets where buyers and sellers negotiate prices, speakers negotiate linguistic features. Feedback occurs as speakers notice what variants of features their communication partners use, and either use the same ones to facilitate understanding or use different ones to create their own identity. Over time, these decisions reinforce each other to create emergent patterns in language – the A-curves that I have been talking about. These A-curve distributional patterns are stable, but because of continuing activity in the system it is always possible that the patterns can change. In human language, we know that they do.

As evidence of this process, let us look at some results from Kretzschmar and Renwick (2016–2018). That project is harvesting vowel measurements automatically from 380 hours of linguistic interviews in the American Southern States.<sup>1</sup> Figures 3 and 4 show two charts of the [æ] vowel (that is, measurements of the vowel in words considered to contain the [æ] vowel in the Carnegie Mellon Pronouncing Dictionary<sup>2</sup>), Figure 3 from a 38-year-old non-African-American woman from Tennessee, and Figure 4 from a 90-year-old African-American man from Alabama. We collected 8939 words from the woman in a three-hour interview, about 2800 tokens per hour. We collected 32424 tokens from the man, about 4000 tokens per hour in a longer interview. We find that the number of tokens we get from different interviews varies greatly owing to the length, circumstances, and recording quality of the interview. The project will harvest vowels from a total of sixty-four speakers, so the effort certainly engages in Big Data, with about two million measurements in the end, as opposed to the [æ] measurements in Figure 1, with just one or two measurements person. Each vowel token was measured at its midpoint, without normalization. The F1/F2 range for the vowel for each person has been divided into a 4×6 grid. The chart counts the number of tokens per grid cell, and then does two things. First, it colors the cells according to a 6-way shading, so that cells with only a few cells are white (cells

---

<sup>1</sup> Methods for extraction of vowel measurements are documented in Miller, Olsen, Renwick, and Kretzschmar (2017). Joey Stanley created the following charts by means of a Shiny app in the R statistical package. The app, called the Gazetteer of Southern Vowels, is accessible at <lap3.libs.uga.edu/u/jstanley/vowelcharts/>.

<sup>2</sup> Available at <www.speech.cs.cmu.edu/cgibin/cmudict>.

with no tokens are not shown at all), and the cells get progressively darker the more tokens they have in them. So, in Figure 3 one cell has the dark color for the most tokens in the 6-way shading, and Figure 4 has two cells with the most tokens (C2 and D3), with two more cells slightly less dense (C3 and D2). The cell counts for the /æ/ vowel for each speaker form an A-curve, below the F1/F2 chart, just what we would expect from the complex system of speech, though the male speaker has four cells with high frequency instead of one or two. Kretzschmar (2015, Chapter 7) shows experimental evidence that distributions with one top-ranked feature occur about 80% of the time, that distributions with two top-ranked features occur about 15% of the time, and that distributions with three or more top-ranked features occur less than 10% of the time: an A-curve of A-curves, what we might expect from the fractal property of speech. The information for these A-curves, of course, is the same information represented with shadings on the chart, using the frequency counts per cell and ordered by the count.

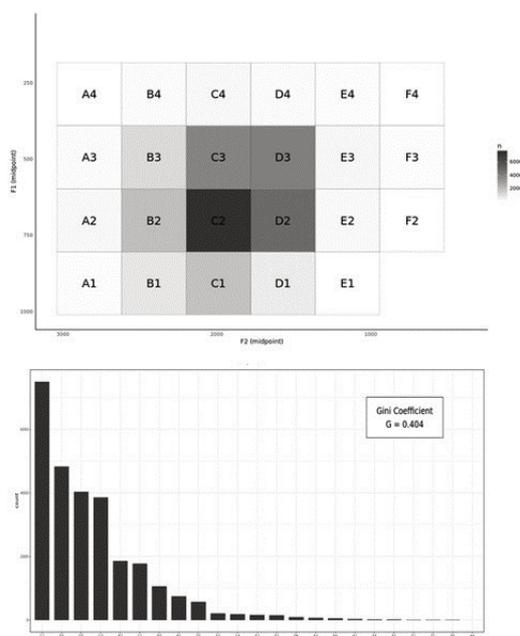


Figure 3. Realizations of [æ] from Speaker 40 (F, W, 38, TN).

The shadings indicate that the dark cells are not exactly in the center of the range, and neither do they appear to be at the center of the distribution. We might rather say that, for this interview, each speaker had a location in F1/F2 space where they most often said their [æ]s, while they had many ways of saying that vowel. This is exactly what complexity science would predict.

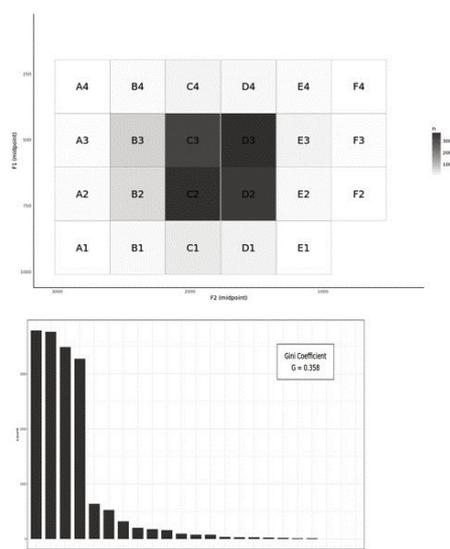


Figure 4. Realizations of [æ] from Speaker 434 (M, B, 90, AL).

The same patterns emerge for the entire data set, not just one speaker, and for just the men in the data set. In Figure 5 we see that the A-curve is still there for all of the [æ] tokens in the entire data set, using 27328 tokens from 47 speakers.

There is one cell, B3, with the highest density of data points. The A-curve below is similar in shape to the A-curve for one speaker. In Figure 6, we see that B3 is again the cell with the darkest shading, when we use 16750 tokens from 24 male speakers (of course this includes part of the same data from Figure 5). The shading pattern for the men, however, is somewhat different from the pattern for all the speakers.

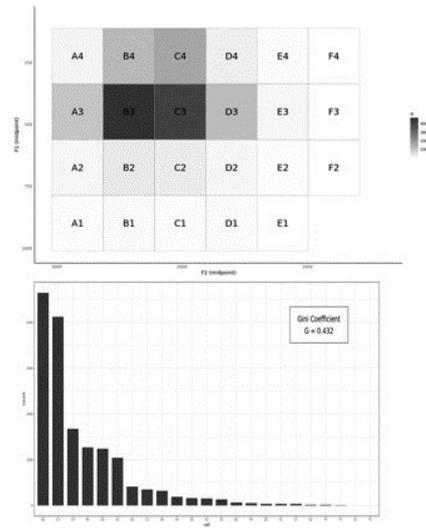


Figure 5. Realizations of [æ] from 47 speakers.

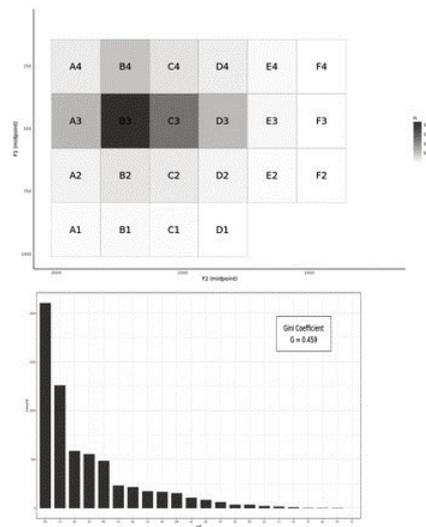


Figure 6. Realizations of [æ] from 24 male speakers.

The “Gini Coefficient” gives a measure of the depth of the curve, and the numbers tell us we have a nonlinear A-curve in both cases. As Kretzschmar (2015: Chapter 7) discusses in detail, normal distributions always have Gini Coefficients below 0.2. The coefficients shown here are all about 0.4 so they all show nonlinear curves, the A-curves I have mentioned.

Here are the tokens we extracted for another vowel, /i/. Again the man in Figure 8 has more tokens than the woman in Figure 7, but both realized the /i/ vowel in a wide range.

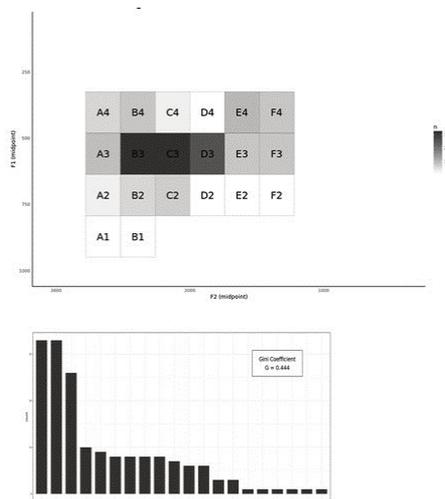


Figure 7. Realizations of [i] from Speaker 40 (F, W, 38, TN).

The woman has two cells in the darkest shading this time (B3 and C3), and so does the man (B4 and C4). The lower charts for each speaker again reproduce an A-curve exactly as complexity science would predict. Again, the shadings indicate that the dark blue cells are not exactly in the center of the range, and neither do they appear to be at the center of the distribution, and we might rather say again that, for this interview, each speaker had a location in F1/F2 space where they most often said their [i]s, while they had many ways of saying that vowel. Again, this is exactly what complexity science would predict.

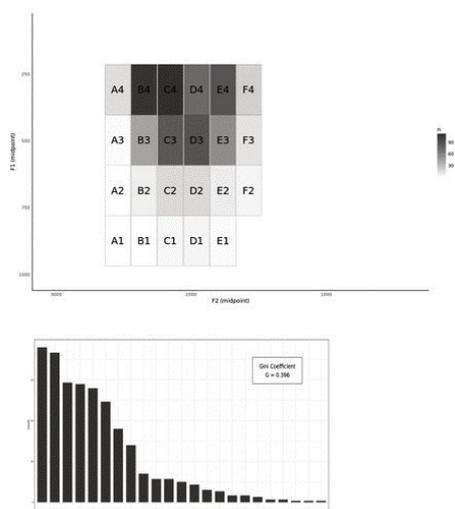


Figure 8. Realizations of [i] from Speaker 434 (M, B, 90, AL).

Figures 9 and 10 show charts of the [i] vowel for the entire data set, and for just the women this time. Again, we see that the shaded boxes in F1/F2 space have one location with darkest shading for the densest pattern of data points, and the A-curves below clearly indicate that the boxes have a nonlinear frequency profile. A close review of the density patterns in F1/F2 space for the single speakers, the men, the women, and the data set overall will show that the patterns and the A-curves for each speaker and for each group of speakers are not exactly the same.

The highest density of data points occurs in a slightly different place on each chart, and the density shadings occur in slightly different places. These differences, as I have said, are how we can recognize each other as individuals and also how we can recognize individuals as members of communities of speakers (whether regional groups or social groups).

In all of the charts there was also wide variability in the vowel measurements, which made for a long tail of cells with low density. Unless we have Big Data and manipulate it like this, we would not be able to see the underlying nonlinear A-curve patterns in a scale-free arrangement. Big Data shows us the homogeneity that Saussure could not see in the totality of language.

Language change works in a new way once we know about the consistent underlying A-curve pattern from the complex system of speech. For

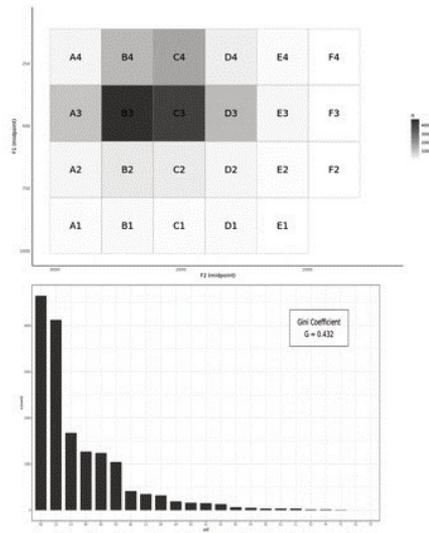


Figure 9. Realizations of [i] from 47 speakers.

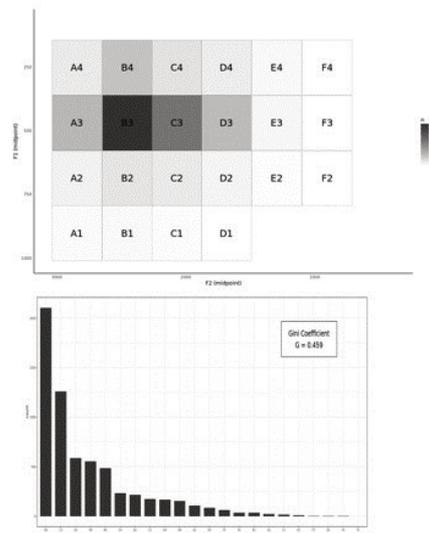


Figure 10. Realizations of [i] from 23 female speakers.

example, new research by Tagliamonte and her colleagues shows that quotative *like* has expanded in four distant cities all at the same time (Figure 11).

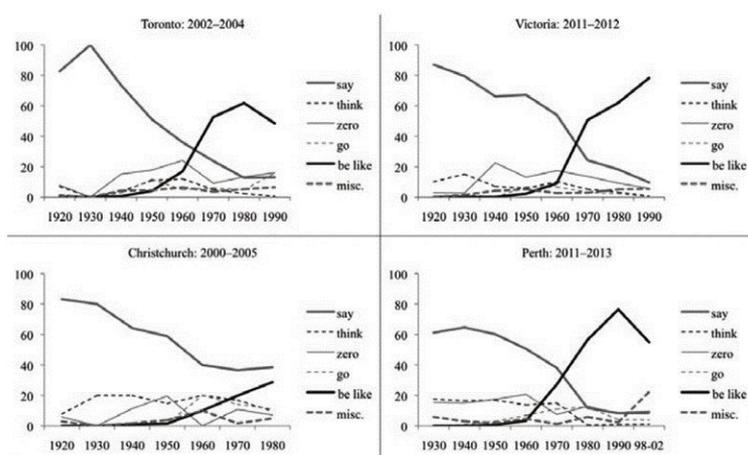


Figure 11. Quotative *like* in four cities (Tagliamonte, D’Arcy, and Rodriguez Louro 2016).

The darkest line shows the growth of quotative *like* in Toronto, Victoria (in western Canada), Christchurch (New Zealand), and Perth (Australia), among people born in the decades indicated at the bottom of each chart. It has surpassed *say* as the quotative marker in all the cities except Christchurch, and even there it is much more frequent among younger people than it is among older people. Other words used as quotatives, like *think*, *go*, and a zero form, behave somewhat differently in each city. However, the large decline of *say* and the large increase of *like* on the A-curve for each place shows us that modern transportation and communication can indeed allow changes in the language to occur on a broad basis across the world, not just in local places. We can make one A-curve to cover the variants in all four of the cities (Figure 12, where five “misc” categories have been represented to replace the single “misc” category shown by Tagliamonte, D’Arcy, and Rodriguez Louro 2016, in order to make the long tail of the curve).

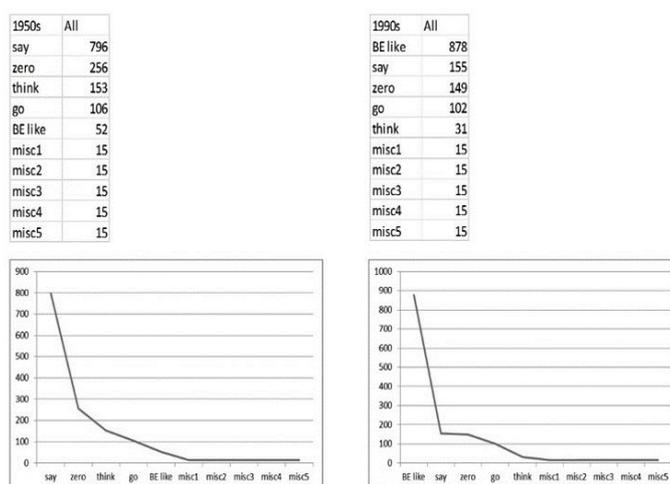


Figure 12. A-curves for quotative variants in four cities, two age groups.

While the frequencies of each of the quotative words have changed between people born in the 1950s and people born in the 1990s, the A-curve is still present for both the older and the younger speakers. Change in the language does not consist of the replacement of one form by another. Change has not meant that the complex system has changed, only that the order of words on the A-curves has changed. This A-curve for all four cities is not identical with the A-curve for any single city – Perth turns out to be closest – but that is what we expect from the complex system: a different A-curve can emerge for each local place, and putting the evidence together at a higher level also produces an A-curve, one that may not match any of them made at lower levels. Of course, this process of observation could be repeated for every feature of English (or any other language), all pronunciations and grammatical usages as well as the choice of word used as a quotative. English can be one language, and yet have different varieties wherever you go.

Saussure was wrong about speech, but of course he did not have the means to be right. Neither did Noam Chomsky, whose influential model of linguistics was developed before we had the means to work with sampling and Big Data. These linguists had good reasons for thinking the way that they did: they needed a way to study language without being able to control

all of the evidence from individual speakers. They created the social fact of linguistic structure, first as a definable generalization based on impressions of the linguist and later *in rebus*, as a real thing. Linguistic structure is definable and graspable as a collective entity because we assume that it is, as Chomsky said, the competence of “an ideal speaker-listener, in a completely homogeneous speech-community, who knows its language perfectly.” Now, however, we know what the evidence of sampling and Big Data shows about speech so we can no longer follow Saussure and Chomsky as far as locating linguistic structure *in rebus*. Once we know that speech has an underlying frequency structure and is fractal in its distributional patterns, we can no longer accept the social fact as the foundation of linguistics. Individuals are all different; speech communities are never homogenous. Instead, speech from different speakers varies continuously, but always in the same frequency profiles, in multiple dimensions in scale-free patterns.

Social facts based on impressions of a linguist can be nothing more than alternative facts, as I described them at the beginning of this essay. Because of the scale-free nature of speech in a complex system, no individual observer can have an idea of the totality of language. Every individual is biased by the particular groups in which she participates, whether regional or social groups, or communities of practice. Every linguist will have a slightly different view of the whole, and could perhaps create a model that was correct for himself but would necessarily be wrong for the unlimited number of groups of speakers of a language. Unless we use sampling methods and Big Data, we are choosing which facts to accept instead of relying on evidence that we can collect. Social facts based on the impressions of a single linguist, or even a group of linguists, can never lead to anything more than alternative facts. Linguists should not just choose what evidence to use when we make generalizations about language change – for example, to say that *like* has replaced *say* in the quotative system – because then our accounts would just be broadcasting fake news. Unless we insist on firm grounding from empirical evidence – sampling and Big Data – linguistics cannot escape from alternative facts and fake news. Those are no better in linguistics than they are in politics.

The way for linguists to improve their generalizations is clear. The new Big Picture for linguistics demands that we begin by identifying a specific community of speakers to study, instead of assuming the overall homogeneity of a language. The new Big Picture for linguistics also demands that our generalizations be based on evidence with appropriate sampling to estimate

the totality of the speech community that we choose to study. Finally, the new Big Picture for linguistics demands that we pay attention to changing frequency profiles, instead of asserting that an entry in the linguistic structure has changed over time. What effect will the new Big Picture for linguistics have upon the research that linguists do? For detailed discussion of that question, I can recommend *Language and Complex Systems* (Kretzschmar 2015). There are chapters on usage-based linguistics, historical linguistics, cognitive linguistics, sociolinguistics, and even the relation of grammars to complex systems. Each of these chapters offers specific suggestions for improvements in these domains of linguistics. In usage-based studies, we cannot expect that the usage we study will automatically fit into a linguistic structure. Grammaticalization, for instance, is just what Paul Hopper said it was in 1987: “[grammatical] structure is always deferred, always in a process but never arriving, and therefore emergent.” His famous later work with Elizabeth Traugott fits the process to a linguistic structure that has arrived; since linguistic structures are static entities, he could do nothing else if, as he has said (personal communication), he was going to accommodate his colleagues. The problem is the accommodation to linguistic structure. We can, however, continue to make grammars if we want to, but the linguistic structures we define will just be snapshots at specific times of the speech of specific populations of speakers, not the generative descriptions of a social fact *in rebus*. If we want to make structural grammars, we should include all of the variants of constructions that we can find. The first question for any construction is not its logical pattern, but how many different ways speakers in a population can implement it. For a generative grammar that focuses on the logical pattern, the syntactician should not try to include all of those different ways to implement a construction, but instead should stick to the most frequent variants, those that are top ranked on the A-curve. This will avoid the problem of rule-creep and highlight the logical relations of frequent constructions. In cognitive linguistics, we need to get past the idea that the brain is like a database that stores symbolic representations, which we then manipulate. I have already given an example from John Miller that illustrates the dynamic process of emergence. The problem with representations in exemplar theory is that they do not address the wide range of variation that actually exists, as for instance in the above illustrations from automatic vowel measurements. The idea of representations is related to the assumption of linguistic structure in usage-based linguistics: if we do not expect linguistic structure *in rebus*, then we are free to prefer Miller’s emergence. Finally, I have already given an ex-

ample, from Tagliamonte and her colleagues, of how we might address historical change and sociolinguistics. The basic problem here is recognition of the fractal nature of speech, that A-curves will exist at all times and in all communities, so that we need to account for change over time as changes in frequency along the A-curve rather than as replacement of features.

It should now be clear that we are at a crucial moment in linguistics. A hundred years ago our most seminal linguists had good reason for preferring linguistic structure, but now we know that their assumptions were incorrect. We have a brand new Big Picture for linguistics, which fits into the scientific model of complex systems. We can embrace our new understanding of language, and in so doing bring new empirical scientific adequacy to our work, or we can live in the past. Any linguist who prefers not to indulge in fake news and alternative facts should make the choice to adopt the complex systems model for human speech.

### References

- Hopper, Pl. 1987. Emergent grammar. *Berkeley Linguistics Society* 13. 139–157. (Available at <<http://home.eserver.org/hopper/emergence.html>>.)
- Kretzschmar, W.A., Jr. 2009. *The linguistics of speech*. Cambridge: Cambridge University Press.
- Kretzschmar, W.A., Jr. 2015. *Language and complex systems*. Cambridge: Cambridge University Press.
- Kretzschmar, W.A., Jr. and M. Renwick. 2016–2018. NSF BCS- 1625680, “Automated Large-Scale Phonetic Analysis: DASS Pilot.”
- Miller, J. 2015. *A crude look at the whole*. New York: Basic Books.
- Miller, R., M. Olsen, M. Renwick and W.A. Kretzschmar, Jr. 2017. Methods for transcription and forced alignment of a legacy speech corpus. *Proceedings of Meetings on Acoustics*, Jul 2017. 1–13.
- Tagliamonte, S., A. D’Arcy and C. Rodriguez Louro. 2016. Outliers, impact, and rationalization in linguistic change. *Language* 92. 824–849.
- Tomasello, M. 2003. *Constructing a language*. Cambridge: Harvard University Press.

### Corresponding author:

William Kretzschmar  
 Department of English  
 Franklin College of Arts and Sciences  
 University of Georgia  
 Athens, GA 30602  
 USA