



April 1996

Mathematical Approaches to Comparative Linguistics

Tandy Warnow
University of Pennsylvania

Follow this and additional works at: http://repository.upenn.edu/ircs_reports

Warnow, Tandy, "Mathematical Approaches to Comparative Linguistics" (1996). *IRCS Technical Reports Series*. 11.
http://repository.upenn.edu/ircs_reports/11

University of Pennsylvania Institute for Research in Cognitive Science Technical Report No. IRCS 96-04

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/ircs_reports/11
For more information, please contact libraryrepository@pobox.upenn.edu.

Mathematical Approaches to Comparative Linguistics

Abstract

The inference of the evolutionary history of a set of languages is a complex problem. Although some languages are known to be related through descent from common ancestral languages, for other languages determining whether such a relationship holds is itself a difficult problem. In this paper we report on new methods, developed by linguists Johanna Nichols (Berkeley), Donald Ringe (Penn), and Ann Taylor (Penn), and computer scientist Tandy Warnow (Penn), for answering some of the most difficult questions in this domain. These methods and the results of the analyses based upon these methods were presented in November 1995 at the Symposium on the Frontiers of Science of the National Academy of Science.

Comments

University of Pennsylvania Institute for Research in Cognitive Science Technical Report No. IRCS 96-04



Institute for Research in Cognitive Science

**Mathematical Approaches to
Comparative Linguistics**

Tandy Warnow

**University of Pennsylvania
3401 Walnut Street, Suite 400C
Philadelphia, PA 19104-6228**

April 1996

**Site of the NSF Science and Technology Center for
Research in Cognitive Science**

IRCS Report 96--04

Mathematical Approaches to Comparative Linguistics

Tandy Warnow*

April 1, 1996

Abstract

The inference of the evolutionary history of a set of languages is a complex problem. Although some languages are known to be related through descent from common ancestral languages, for other languages determining whether such a relationship holds is itself a difficult problem. In this paper we report on new methods, developed by linguists Johanna Nichols (Berkeley), Donald Ringe (Penn), and Ann Taylor (Penn), and computer scientist Tandy Warnow (Penn), for answering some of the most difficult questions in this domain. These methods and the results of the analyses based upon these methods were presented in November 1995 at the Symposium on the Frontiers of Science of the National Academy of Science.

1 Evolutionary relationships in linguistics

Evolutionary relatedness of languages is described by observing that separation of speech communities into separate and noninteracting sub-communities eventually results in a language developing into distinct new languages in a process quite similar to speciation in Biology. While this is not the only means by which languages change, it is this process which is referred to when we say, for example, “Latin is an ancestor of French.” This allows us to model the evolution of related languages as a rooted tree in which internal nodes represent the ancestral languages. When a set of languages does not have a common ancestor (such as may be the case for a set containing both Dravidian and Indo-European languages), then the evolution of that set is best described by a disjoint collection of rooted trees (i.e. a *forest*). Except in circumstances involving related dialects

*Department of Computer and Information Science, University of Pennsylvania Philadelphia, PA. Email: tandy@central.cis.upenn.edu. This work was supported in part by NSF grants CCR-9457800 and SBR-9512092, by an ARO grant DAAL0389-C0031, by the Institute for Research in Cognitive Science, and by financial support from Paul Angello.

which continue to have close contact, there is no problem with this model of language evolution.

Careful scholarship over the last century has determined critical features and patterns that, combined with a statistical analysis, can be used to establish that languages share a common ancestor; examples of these features are shared idiosyncracies in the grammars, shared idiosyncratic sound changes, and patterns of sound correspondences. Extending this fundamental statistical analysis, two techniques (the *comparative method* and *subgrouping through shared innovations*) have been developed which enable linguists to infer greater information about relatedness and properties of ancestral languages, and - to a limited extent - subgrouping as well. These techniques have established all known linguistic families and subfamilies, and are the basis of historical linguistic scholarship. Known families presently number close to 300, though ongoing comparative work on the languages of New Guinea and of South America - two of the linguistically most diverse and least described places on earth - may reduce this total to as low as 200. Many of these "families" are in fact one-descendent ones, like Basque, which is a distinct genetic lineage of its own with no known kin. Although these techniques provide firm evidence of relatedness between languages, they provide only limited information (as previously applied) about subgrouping within sets of related languages, and thus do not provide a method by which linguists can infer the full evolutionary history of language families. Consequently, the evolutionary histories of many language families have been debated for decades without resolution.

Finally, these techniques are only applicable for comparing well-attested languages which are known to be related and whose most recent common ancestor does not lie more than 6,000 - 8,000 years in the past. At time depths beyond that limit, the critical features upon which the classical techniques are based survive in such small numbers that they cannot reliably be distinguished from chance resemblances[23]. Attempts have been made to establish criteria by which such relationships can be inferred for sets of languages with ancestors further back in time than this barrier, but these have been largely unsuccessful and heavily criticized for lacking rigorous statistical foundations. Extending the range of linguistic comparison beyond that critical time depth is therefore a major endeavor within historical linguistics.

In the Frontiers of Science symposium, the panel on Mathematical Approaches to Comparative Linguistics discussed new approaches to solving these two major problems. The first talk involved a team of two linguists, Donald Ringe and Ann Taylor, and a computer scientist, Tandy Warnow, in their efforts to develop a methodology for inferring the evolutionary tree for languages known to be related. They formulated a model of evolution based upon classical scholarship in historical linguistics, and developed an efficient method which would serve two purposes: first, the model could be tested to see if it fit the data, and trees which best fit the model could be generated. The application of their methods to the Indo-European family of languages has indicated that the

data to a great extent fit the model extremely well, and produced a robust evolutionary tree, potentially settling longstanding controversies in Indo-European studies. In the second talk, Johanna Nichols of UC Berkeley described her method by which relationships and/or earlier interaction could be reliably inferred between languages not necessarily known to be genealogically related. She described properties of linguistic features which she called *population markers* which would reliably indicate either genealogical relationship or at least significant and prolonged contact between language communities. Her analysis of the world’s languages has implications for our understanding of human migrations and greatly extends the power of comparative linguistic analysis.

In this report, we will describe the basic ideas and results of these two research projects, and report on some of the questions posed by members of the audience at the Symposium. Each of these projects is ongoing, with developing methodologies and continuing data analyses. Consequently, some of the results are new and did not appear in the Symposium.

2 Ringe, Warnow, and Taylor

The two fundamental techniques for subgrouping within established families used in Historical Linguistics are the Comparative Method, formalized by Henry Hoenigswald in [14], and subgrouping through shared innovations. Since the assumptions upon which these two techniques are based are used in the methodology developed by Ringe, Warnow, and Taylor, we describe these techniques in some detail.

The Comparative Method: Given a set of languages known to be related, the comparative method has the following steps: *Step 1: Observe sound correspondences*; that is, compare words for the same (or comparable) meanings, and observe patterns of sound correspondences between pairs of languages. *Step 2: Infer regular sound change rules*. These rules must explain *all* the sound correspondences observed in Step 1. These rules may be context free or context dependent, and are specific to each lineage. *Step 3: Infer cognation judgements*. Two words w and w' from two languages L and L' respectively are said to be *cognate* if it is possible to infer a word w^* in some common ancestor of L and L' such that each w and w' can be derived from w^* by the sound change rules specific to L and L' , respectively. The comparative method distinguishes between words that are similar and those which have a common origin, and thus enables linguists to establish that Spanish **mucho** and English **much** are not cognate, because applications of the sound change rules do not indicate that they come from a common ancestral word (**mucho** is derived from **multum** in Latin, meaning “much”, while **much** is derived from **micel** in Old English, meaning “big”).

Linguistic characters The comparative method defines cognate classes so that different words may be considered to be *equivalent*, and thus allows the languages to be defined by a set of equivalence relations, one for each meaning. This is comparable to using morphological features or columns within biomolecular sequences to represent biological taxa; in each case, the primary data are described through the use of partitions of the taxa into equivalence classes. Such partitions are called *characters* in the biological literature.

The comparative method establishes two types of linguistic characters, *lexical* and *phonological*. For lexical characters, the character is the semantic slot, as for example, the meaning ‘hand’, with the states of the character defined by cognation judgements. (Were it not for *word replacement*, which is endemic across all languages, words for the same meaning in related languages would all be cognate, and thus all lexical characters would have a single state on any set of related languages. Thus, word replacement is the reason that lexical characters have more than one state.) For phonological characters, the character is a sound change. Languages which share the same outcome (generally, those that undergo the change versus those that do not) exhibit the same state for the character. As a special subtype of lexical characters, *morphological* characters can also be defined. Here, the character is generally a grammatical feature, for example, the formation of the future stem, the way the passive is marked, the genitive singular ending of o-stem nouns and adjectives, etc. Languages in which the feature is instantiated in the same way, or by a reflex of the same proto-morpheme, exhibit the same state for the character. Because morphological characters resist borrowing, they are especially useful in determining relationships between languages.

Subgrouping through shared innovations: Classical methodology in historical linguistics has used these phonological and morphological characters for subgrouping purposes by noting that when a character has two states in which one is clearly ancestral, then the character defines a linguistic innovation. Linguistic innovations which are useful for subgrouping must be peculiar enough to not be easily repeated, and (depending upon the particular set of languages being examined) should not be too easily lost. When a statistically significant number and quality of innovations are shared, then the set of languages sharing that common set of innovations can be considered to form a linguistic *subgroup*, such as the Germanic and Italic subfamilies of Indo-European.

Comments The classical methodology in historical linguistics is surprisingly powerful. As we have shown, cognation judgements derived from rigorous application of the comparative method are not measures of similarity (otherwise **mucho** and **much** *would* be cognate) but of **homology** (descent from a common origin). Furthermore, when languages are very well attested, the comparative method enables linguists to detect almost all instances of borrowing; thus, the

application of the method implies that all words in English beginning with **sk** are borrowed from other languages (for example, **sky** is borrowed from Old Norse and **skunk** is borrowed from Algonkian).

There are some limitations to these classical techniques, however. Word replacement is such a relatively frequent phenomenon that after a period of approximately 6,000 years, it is essentially impossible to detect cognates; other diagnostic features of languages are also gradually lost, and thus the detection of relatedness between languages is a difficult task at large time depths. In addition, these classical methods require that the languages be well attested (so that, for example, the sound change rules can be complete and accurate); thus, even for closely related languages (i.e. those with common ancestors that are not too far back in time), inferring the subgrouping within the family, or even the relatedness of such languages, can at times be difficult.

Despite these limitations, classical methodology has successfully identified the major families and subfamilies (Germanic, Indo-European, Dravidian, etc.) of the world's languages. The reason these methods have not successfully resolved controversies about subgrouping within established families is that the method for subgrouping has required very restrictive properties about the data used for that purpose. Thus, these methods have been more useful for recognizing relatedness rather than subgrouping purposes.

The key observation made by Ringe and Warnow in the fall of 1993 that enabled them to develop a new methodology was that the classical methods in Historical Linguistics (subgrouping through shared innovations and the Comparative Method) can be stated as hypothesizing that *almost all linguistic characters, if properly encoded, should be compatible with the evolutionary tree for the languages*. The term *compatible* is a technical term from the systematic biology literature, which has the following definition: *a character c is compatible with tree T if the nodes in T can be labelled by states of c so that every state of c induces a connected subset of T* . An example of a biological character which is compatible is the *vertebrate-invertebrate* character, while the character indicating the presence or absence of wings is not a compatible character on the tree of all animals.

The reason that the hypothesis is stated with the caveat that only *almost all* and not *absolutely all* characters should be compatible is the observation that many phonological characters are based upon sound changes that are natural enough to occur repeatedly. By contrast, lexical characters ought to be compatible on the evolutionary tree, provided that borrowing can be detected. Those morphological characters and phonological characters that are based upon properties unusual enough to have only arisen once also ought to be compatible on the evolutionary tree. Thus, the hypothesis indicated by the classical methodology is, more precisely, that all lexical characters, and those morphological and phonological characters which represent distinctly unusual traits, should be compatible on the evolutionary tree of a family, provided that the family is well attested and well understood.

Although the linguistic hypothesis is that all properly selected and encoded characters *should* be compatible on the true evolutionary tree, there are certain specific conditions in which it can be difficult to distinguish between true cognates and words which are borrowed; that is, it may be difficult to distinguish between true and false cognates. Based upon these observations, Ringe and Warnow formulated the following optimization criterion: *find the tree on which it is possible to explain all incompatible character evolution with as simple an explanation as possible, and which matches linguistic scholarship as closely as possible.*

The optimization problem they formulated is related to a classical problem in biological systematics called the *Compatibility Criterion*, in which the tree on which as many characters as possible are compatible is the optimal tree. The compatibility criterion problem caught the interest of the computer science algorithms community because of its combinatorial flavor and interesting graph-theoretic formulation[6]. In addition to showing that the compatibility criterion problem is NP-hard [4, 10, 25] (and thus unlikely to be solvable in polynomial time - see [12]), computer scientists and mathematicians developed polynomial time algorithms for various fixed-parameter formulations of the problem[15, 16, 17, 1, 2, 3]. Using a program designed by Richa Agarwala (based upon [2]) to solve the compatibility criterion, Warnow and Ringe decided to test the hypothesis of classical historical linguistics that properly encoded linguistic data should result in highly compatible characters. The program in turn would also permit them to explore all the trees which had optimal and near-optimal scores for the compatibility criterion, and thus select those trees with (hopefully) simple explanations of incompatibility.

Assisted by Libby Levison, then a doctoral candidate at Penn, Ringe and Warnow first tested this hypothesis on some small data sets. These preliminary results were very encouraging, and Ringe and Warnow then turned to the Indo-European (IE) family. Although the IE family is among the best understood of the world's language families, there are several longstanding controversies in the field, specifically the Indo-Hittite and Italo-Celtic hypotheses, which had resisted analysis. (The Indo-Hittite hypothesis is that the first subfamily to break off from the root of the Indo-European evolutionary tree should be the Anatolian branch, represented by Hittite, and the Italo-Celtic hypothesis is that Italic and Celtic should be sisters within the tree, and without a third sister.) Ringe and Warnow were interested in seeing whether their techniques would help resolve any of these difficult issues.

They selected from each of the subfamilies within IE the oldest well-attested language to represent the subfamily. In order to reduce the possibility of borrowings among the lexical characters and bias on their part in choosing these characters, they used an existing basic vocabulary list of 208 semantic slots[27].¹

¹The list has one more item than Tischler's[27] because we split the item *day* into two items, *period of 24 hours* and *period of daylight*.

Each semantic slot was treated as a single character and judgements of cognation were made on the basis of the comparative method. Since nothing similar to a basic vocabulary list exists for morphological and phonological characters and since these will vary from family to family, an appropriate set of morpho/phonological characters has to be developed for each family. For the IE test they used ten Proto-Indo-European morphological items which have a reflex in most of the IE languages, and four phonological developments which they judged to be sufficiently abnormal as not to be easily repeatable.

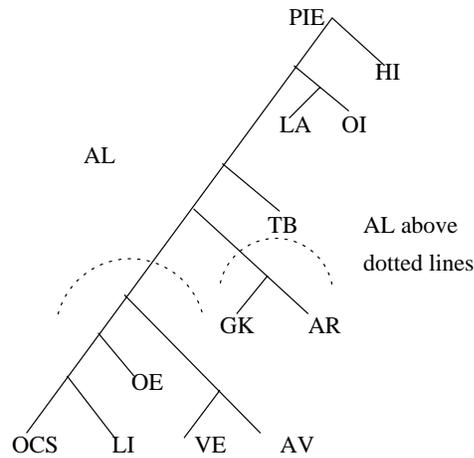
Over the next two years, in collaboration with postdoctoral researcher Ann Taylor, Ringe and Warnow studied the Indo-European family of languages. They discovered that a phenomenon they termed *polymorphism* creates significant difficulties for reconstructing the evolutionary history in Indo-European. Since no methodology had yet been established for handling polymorphic data, Warnow (in collaboration with other computer scientists) developed algorithms to handle polymorphic character data[5]. Because rooted trees are desirable, directionality constraints implied by some of the linguistic data were encoded as characters, using techniques already in use by systematic biologists, and these characters were included in the dataset.

These algorithms were then applied to the entire data set for Indo-European, and all the trees with optimal or near-optimal compatibility scores were examined. The data included 222 characters describing 12 languages. The two best trees had 12 and 13 incompatible characters respectively, but the rest were significantly worse. Although this is a high number of incompatible characters, the ultimate analysis of the linguistic characters on the dataset when Germanic was removed supported a *perfect phylogeny* (i.e. all the characters were compatible on that tree). (A perfect phylogeny is the best possible tree with respect to both compatibility and parsimony criteria.) Furthermore, the pattern of incompatibility has a simple explanation: it appears to point to a situation in which Germanic began to develop within the Satem Core (as evidenced by its morphology) but moved away before the final satem innovations. It then moved into close contact with the “western” languages (Celtic and Italic) and borrowed much of its distinctive vocabulary from them at a period early enough that these borrowings cannot be distinguished from true cognates. Because statements of cognation depend upon unbroken descent from a common ancestor through *genetic* inheritance, and not from borrowing, this hypothesis implies that words in Germanic borrowed from pre-*proto-Italic* and pre-*proto-Celtic* are not cognate with the corresponding words in Italic and Celtic. Thus, this hypothesis, simple as it is, results in a revision of the encoding so that all the characters are compatible on the tree.

A careful examination of set of optimal and near optimal trees shows that the near-optimal trees all differ from the optimal tree (shown in Figure 1) in only minor ways. The root (PIE, indicating Proto-Indo-European) can move down the tree somewhat, the placement of Tocharian can move out of the branch containing Greek and Armenian to just above it, and the edge separating the

pair Italic-Celtic from the rest of the languages can be contracted. Some of the differences (such as the location of the root) can be resolved on the basis of geographical feasibility. Thus, the analysis indicated consistent support for the Indo-Hittite hypothesis, surprising Ringe who had frequently argued against this hypothesis, and weak support for the Italo-Celtic hypothesis, which currently rests upon exactly one character.

The authors concluded with noting that this method permits a linguist to determine the tree or trees which best fit his or her interpretations of the linguistic data, and thus to test the consistency of the linguist's judgements.



HI - Hittite, OI - Old Irish, LA - Latin
 TB- Tocharian B, GK - Greek, AR - Armenian
 OE - Old English, OCS - Old Church Slavic
 LI - Lithuanian, AV - Avestan, VE - Vedic
 AL - Albanian

Figure 1: The tree on the Indo-European data set.

Questions Some of the questions posed by the audience enabled the team to clarify their methods and findings, so that this report hopefully provides the answers. These questions were: (1) *Why is a tree the correct model of linguistic evolution? What about creoles and pidgins?* Answer: because language communities separating defines a rooted tree, and mixed languages such as creoles and pidgins can be detected as such, and do not cause problems for the inference of evolutionary history. (2) *Why is compatibility the right optimization criterion?* Answer: we're using the compatibility criterion as a way of *testing*

the assumptions upon which classical methodology is based. (3) *Why do you believe this tree is really the correct tree?* Answer: although we have a fair degree of confidence that this is the correct tree, the only thing we are sure of is that it is the tree that best fits the assumptions of the linguistic scholarship and our interpretations of the data. What we do have great confidence in are those features that remain constant across the set of all the near-optimal trees.² (4) Noting that data used in the study did not include some phonological characters, because these characters were based upon sound changes that were too easily repeated (such as the loss of the initial **h** in words), one member of the audience asked whether this wasn't potentially cheating, eliminating characters that simply didn't fit our pre-conceived notion of what the correct tree was. To this, the authors replied that such judgements seem inevitable. Comparable judgements arise in the analysis of morphological data in Biology, where characters such as *presence or absence of a backbone* and *presence or absence of wings* cannot be treated identically. On the other hand, the authors noted that all lexical characters (i.e. those based upon cognation judgements) and morphological characters were included in the final analysis; only phonological characters based upon natural sound changes which are easily repeated were removed from the data set. Thus, Ringe, Warnow, and Taylor felt that extremely high compatibility scores that resulted from their analysis indicated that the hypothesis they tested (that linguistic characters are compatible on the evolutionary tree) seems to be valid to a large degree for the Indo-European family, and thus the evolutionary trees with high compatibility scores are potentially the best candidates for being the true evolutionary tree.

3 Johanna Nichols

The previous section described how the evolution of a set of languages sharing a common origin can be inferred from the features of the languages when properly encoded and analyzed. Johanna Nichols' work studies the case of languages which are either unrelated, or which have diverged to the point where the diagnostic features used to infer genetic relatedness between languages have been largely lost. Although researchers from various fields have attempted to establish techniques by which genetic relationships can be reliably inferred between distant languages, such techniques have been largely unsuccessful and heavily criticized within the historical linguistics community for their infirm statistical foundation. One of the reasons this endeavor is particularly difficult is that after periods of approximately 6,000 to 8,000 years, it is difficult to distinguish between similarities due to common origin and those due to prolonged and in-

²Reconsideration of the data and gathering and analyzing more data, are part of the research effort, and resulted – after the presentation at the Symposium – in a slightly revised hypothesis of the evolutionary history for Indo-European. Thus, the tree that we present here is slightly different from that presented at the Symposium.

timate contact between speech communities. Recognizing this, Nichols' work endeavors to establish techniques by which similarity due to common origin *or* prolonged and intimate contact can be established. She proposes specific features, which she calls *population markers* or *historical markers*, whose distribution can be used to formulate hypotheses about linguistic prehistory. Nichols suggests that her results can be used in conjunction with archeological evidence to develop better theories about early human migrations. Her findings, applied to a database of the world's populations, have the potential to greatly extend current knowledge of human migrations and relationships between languages.

Genetic vs. historical markers *Genetic markers* are features that indicate a genetic relationship between languages, and thus indicate that languages sharing the genetic marker have a common ancestor. By contrast, *historical markers* (also called *population markers*) indicate a non-accidental relationship, though they cannot tell us whether that relationship is specifically genetic; it could have been significant prior contact between speech communities, or prior contact with a now-defunct third party. There are essentially three mechanisms by which languages can share features:

1. through inheritance from a common ancestor, indicating a genetic relationship,
2. through borrowing (whether direct or indirect) between neighboring speech communities, indicating a historical (but not necessarily genetic) relationship, and
3. through spontaneous reappearance of the same trait in different lineages.

In order for any feature to be useful for detecting genetic or historical relationship, the feature must be unlikely to evolve spontaneously; otherwise, spurious relationships will be posited. To establish a specifically genetic relationship (as opposed to the more general historical relationship), it must be possible for the linguist to distinguish between acquisition through borrowing and acquisition through inheritance. Features which are difficult to borrow are appropriate for use as genetic markers, but note that borrowable features can be analyzed correctly in genetic terms, provided that borrowing can be detected, and can thus be useful as genetic markers. Essentially, therefore, genetic markers must have the following traits:

1. The feature must be extremely unlikely to arise twice; for lexical characters, the comparative method establishes this strong probability, and
2. Borrowing of the feature must either be extremely unlikely, or it must be possible to detect such borrowing.

As Ringe and Warnow observed (and subsequent research with Taylor supported), it follows that *genetic markers should define characters which are compatible on the evolutionary tree for the language family*. This observation allows a linguist to posit that some set of features is inherited genetically, and this hypothesis in turn can be tested (using the methodology of Ringe, Warnow, and Taylor) described in the previous section.

Historical markers must also have certain properties that enable a historical relationship to be detected, although these properties are somewhat different from those required for genetic markers. Although the trait should not be likely to arise twice, the condition that borrowing should either be unlikely or detectable need not hold. If a historical marker is based upon a trait which is never borrowed, then it cannot be used to provide evidence of contact between different languages otherwise not known to be related. On the other hand, if the trait is too easily borrowed, or too easily lost, then there will be no pattern of relationship that permits nontrivial observations. Thus, historical markers, to be useful, must be capable of being borrowed, but must not be lost too easily once acquired.

Each type of marker (genetic or historical) enables the detection of a relationship of some sort, either through descent from a common ancestor or through contact, and the best markers (whether genetic or historical) are low-frequency features that form a single frequency peak or cluster, resulting in a frequency asymmetry that is statistically significant. Genetic markers such as these permit subgrouping at a fine-grained level, while historical markers of this type provide greater insight into the history of early human migrations, because the findings can be compared to archeological evidence.

Nichols' proposes a method by which historical markers can be selected and analyzed. She shows how the geographical distribution of a candidate historical marker among the world's languages can provide evidence for common histories between languages, and in particular can lead to hypotheses about early migrations which can then be tested against archeological evidence.

Nichols' research Nichols selected fourteen (14) different traits which had the specified properties required for historical markers, and which in addition were believed to be independent of each other. These were *morphological ergativity, morphological complexity, head-marking morphology, inclusive/exclusive oppositions in first person pronouns, genders or other noun classes, numeral classifiers, tones, possessive affixes, regular transitivization in verbal derivation, identical stems in "I/me" and "we/us", m as root consonant in first person singular pronoun ('I/me'), m as root consonant in second person singular pronoun ("you"), verb-initial word order, and secondary glottal articulation*.

A selection of the world's languages was then studied to determine the incidence of these traits throughout the world. Historical linguistics has established approximately 200-300 different language families, with some of these families

very well understood and others less so. Because some languages are only recently attested (and not as well studied as others), there is a distinct possibility that in time, linguistic scholarship will be able to identify genetic relationships between certain families. Thus, the number of linguistic families may in time be reduced to about 200; that is, languages that now seem to be unrelated genetically, may in time be established as having a common origin. In developing a database of the world's languages, Nichols selected at most one language from each major branch within each linguistic family, to obtain her sample of languages. The sample she has obtained (of over 200 languages, and still growing) has the property that no two languages within the sample are likely to be more closely related than two distantly related Indo-European languages (like French and Armenian).

Geographical distribution of markers Nichols discovered striking patterns in the geographic distribution of these historical markers around the world. All findings point in the same direction: strong affinities between Australia and the western Old World and different but also strong affinities all around the Pacific Rim. The linguistic distributions point to coastal spread around the Pacific beginning in very early times and to an earlier expansion from Africa via southern Asia to Australasia. Both expansions are widely assumed by archeologists and human geneticists, but the linguistic distributions seem to provide the clearest evidence of them.

For example, some markers are most frequent in Europe, Africa, or both, least frequent in Australia, and of middling frequency in Asia and the Americas. This geographical distribution correlates with archeological research that establishes that the Americas were settled by people migrating from Siberia (i.e. from Asia). Other markers are densely clustered in Australia, well represented around the Pacific and in the Americas, but rare in the Old World (Europe, Africa, central Asia), implying that the distribution of these markers must have taken place before the colonization of the Pacific Islands and the New World. The pattern also suggests that the impetus for expansion came from the west, ultimately from Africa.

A similar pattern occurs within Australia and New Guinea, where the frequencies of population markers show that a subset of the Australian languages defined by specific geographic boundaries closely resembles a subset of the languages of New Guinea, again defined by specific geographic boundaries. Other interesting correlations between Australia and New Guinea show up in this analysis, showing generally an east-west trend in the frequencies of the different markers. It is known that Australia and New Guinea were originally (during the Ice Age) parts of the same continent which was split by a postglacial sea-level rise. It is also known that human colonization of these two lands emanated from Southeast Asia, and that the landfall point for this colonization was the northwest coast of the continent. The patterns between these two lands actu-

ally indicate multiple linguistic colonizations, and support the previous research indicating that human colonization occurred when the two lands were in a single continent.

Thus there are many striking patterns that can be observed when the frequencies of these population markers are compared with geography, and these patterns, when combined with archeological evidence, provide significantly more detailed information (or at least better hypotheses!) about early migrations.

Questions The questions posed by the audience mostly focused on issues regarding how linguists determine how far apart languages are: *do linguists assume a constant rate of change in languages?* (answer: not really); *do such assumptions matter?*, (answer: no, not for these purposes, as has been shown in the analysis); and (more generally), *how do linguists decide what a language is as opposed to a dialect?* (answer: “a language is a dialect with an army!”; more seriously, languages are mutually unintelligible forms of speech). They were also concerned with issues regarding the selection of population markers, since predispositions based upon studies for one language family can affect the final conclusions of the research. To this, Nichols responded that she has drawn all usable features from the typological literature and consulted with specialists on various language families, and was recruiting into this kind of work specialists on different language families, and this would help avoid the introduction of bias into the sample. The final questions were whether it was generally believed that all languages are genetically related, whether language (or speech) arose once, and how such questions might be addressed. To this, Johanna Nichols’ answer was essentially that there is no consensus on these questions, and potentially no way of answering them; however, the migration patterns suggested by markers, though very ancient, are much younger than the rise of modern humans and thus, perhaps, younger than the rise of human language.

4 Further Reading

A good introduction to phylogenetic tree construction methodology in Biology can be found in [11]. The methodology of the Ringe, Warnow, and Taylor research is described in [5, 29]. More detailed information about the mathematics of the compatibility criterion problem can be obtained in [16, 28]. Additional material on historical linguistic methodology can be obtained in [13, 14, 18]. Johanna Nichols’ work is described in greater detail in [19, 20, 21]. Discussions of the Indo-Hittite and Italo-Celtic hypotheses (and other controversies in Indo-European studies) can be found in [7, 8, 9, 22, 26]. A discussion of the archeological evidence related to the discussion of migrations in Australia and New Guinea can be found in [30, 24]. For further information on these research projects, the authors may be reached by email:

tandy@central.cis.upenn.edu, dringe@unagi.cis.upenn.edu, ataylor@linc.cis.upenn.edu,
and johanna@uclink.berkeley.edu.

5 Bibliography

References

- [1] R. Agarwala and D. Fernández-Baca. *A Polynomial-Time Algorithm for the Perfect Phylogeny Problem when the Number of Character States is Fixed*. SIAM Journal on Computing, Vol. 23, No. 6, pp. 1216–1224, December 1994.
- [2] R. Agarwala and D. Fernández-Baca. *Fast and Simple Algorithms for Perfect Phylogeny and Triangulating Colored Graphs*. To appear in the special issue on *Algorithmic Aspects of Computational Biology* of International Journal of Foundations of Computer Science. Available as DIMACS technical report TR94-51.
- [3] R. Agarwala, D. Fernández-Baca, and G. Slutzki. *Fast Algorithms for Inferring Evolutionary Trees*. Journal of Computational Biology, Vol. 2, No. 3, pp. 397–408, Fall 1995.
- [4] Bodlaender, H., Fellows, M. and Warnow, T. 1992: Two strikes against perfect phylogeny, Proceedings of the International Congress on Automata and Language Processing.
- [5] Bonet, M., Phillips, C., Warnow, T. and Yooseph, S. 1995, *Constructing evolutionary trees in the presence of polymorphic characters*, to appear, Proceedings of the 1996 ACM Symposium on the Theory of Computing, Philadelphia PA.
- [6] P. Buneman, *A characterization of rigid circuit graphs*, Discrete Mathematics, 9 (1974), 205-212.
- [7] Cowgill, Warren 1970: Italic and Celtic superlatives and the dialects of Indo-European, in Cardona, George, Henry M. Hoenigswald, and Alfred Senn (eds.), *Indo-European and Indo-Europeans*, University of Pennsylvania Press, Philadelphia.
- [8] Cowgill, Warren 1975: More evidence for Indo-Hittite: the tense-aspect systems, in Heilmann, Luigi (ed.), *Proceedings of the Eleventh International Congress of Linguists*, Mulino, Bologna.
- [9] Cowgill, Warren 1979: Anatolian *hi*-conjugation and Indo-European perfect: instalment II, in Neu, Erich, and Wolfgang Meid (eds.), *Hethitisch und Indogermanisch*, Innsbrucker Beiträge zur Sprachwissenschaft, Innsbruck.

- [10] W. H. E. DAY AND D. SANKOFF, *Computational complexity of inferring phylogenies by compatibility*, Syst. Zool., Vol. 35, No. 2 (1986), pp. 224–229.
- [11] Felsenstein, J. 1982: Numerical methods for inferring evolutionary trees, *The Quarterly Review of biology*, Vol.57, No.4.
- [12] Garey D. and Johnson, D: *Computers and Intractability*, Freeman, New York, 1979.
- [13] Greenberg, JH, 1957: Genetic relationship among languages. In Greenberg, JH (ed), *Essays in Linguistics*, pp. 35-45, Chicago, University of Chicago Press.
- [14] Hoenigswald, Henry M. 1960: *Language Change and Linguistic Reconstruction*, University of Chicago Press, Chicago.
- [15] Kannan, S. and Warnow, T. 1994: Inferring evolutionary history from DNA sequences, *SIAM Journal on Computing* 23(4):713-737.
- [16] Kannan, S. and Warnow, T. 1995: A fast algorithm for the computation and enumeration of perfect phylogenies, to appear, *SIAM Journal on Computing*.
- [17] McMorris, F.R., Warnow, T. and Wimer, T. 1994: Triangulating vertex colored graphs, *SIAM Journal on Discrete Mathematics* 7(2):296-306.
- [18] Meillet, A. 1925: *La Méthode Comparative en Linguistique Historique*, H. Aschehoug & Co., Oslo.
- [19] Nichols, J. 1990: Linguistic diversity and the first settlement of the New World. *Language* 66: 475-521.
- [20] Nichols, J. 1992: *Linguistic Diversity in Space and Time*, Chicago: University of Chicago Press.
- [21] Nichols, J. Diachronically stable structural features. In Anderson H (ed) *Proceedings of the 11th International Congress of Historical Linguists*.
- [22] Ringe, Donald A., Jr. 1991: Evidence for the position of Tocharian in the Indo-European family? *Die Sprache* 34:59-123.
- [23] Ringe, Donald A., Jr. 1992: On Calculating the Factor of Chance in Language Comparison, *Transactions of American Philosophical Society* Vol.82, no.1, Philadelphia, PA.
- [24] Roberts RG, Jones R, Smith MA (1990): *Thermoluminescence dating of a 50,000-year-old human occupation site in northern Australia*, *Nature* 345: 153-156.

- [25] Steel, M. 1992: The complexity of reconstructing trees from qualitative characters and subtrees, *Journal of Classification* 9:91-116.
- [26] Sturtevant, Edgar H. 1933: *A comparative grammar of the Hittite language*. Linguistic Society of America, Philadelphia.
- [27] Tischler, Johann 1973: *Glottochronologie und Lexikostatistik*. Innsbrucker Beiträge zur Sprachwissenschaft, Innsbruck.
- [28] T. Warnow, *Efficient algorithms for the character compatibility problem*, New Zealand Journal of Botany, Vol. 31 (1993), pp. 239-248.
- [29] Warnow, T., Ringe, D., and Taylor, A. *Reconstructing the evolutionary history of natural languages*, Proceedings of the ACM/SIAM Symposium on Discrete Algorithms, 1996.
- [30] White JP, O'Connell JF 1982: *A Prehistory of Australia, New Guinea, and Sahul*, New York: Academic Press.