

Generalizing Scales

Michael Cysouw

Abstract

Instead of considering scales to be linearly ordered structures, this paper proposes that the linguistic notion of a scale can be fruitfully generalized as a special case of a metric. Further, to be considered a scale of typological interest, there should be a significant correlation between a meaning-scale and a form-scale. As a hands-on example of the proposals put forward in this paper, the “scale of likelihood of spontaneous occurrence” (Haspel-math 1993) is reanalyzed. This scale describes the prototypical agentivity of the subject of a predicate.

1. Scales as restrictions on form-function mapping

Scales¹ of linguistic structure are one of the more promising avenues of research into the unification of the worldwide linguistic diversity. Although our growing understanding of the diversity of the world’s languages seems to throw more and more doubt on the many grandiose attempts at universally valid generalizations, the significance of scales (such as the well-known animacy scale) for human languages still appears to stand strong (for a different opinion on this specific example see Bickel & Witzlack-Makarevich, this volume).

But what exactly is a scale? A scale seems to be mostly thought of as an asymmetrical one-dimensional arrangement (a “total order” in mathematical parlance) of certain cross-linguistic categories/functions. Put differently, a scale is a linear ordering of functions with a “high end” and a “low end”. To be considered an interesting scale, the formal encoding of these functions in actual languages should be related to this linear ordering. For example, specific encodings should typically be restricted to either end of the scale, like nominative and ergative encoding on the animacy scale (Silverstein 1976).

In this paper, I will argue that this concept of a scale can be fruitfully generalized. In a very general sense, all linguistic structure consists of forms expressing particular functions. If we find restrictions across languages on the kind of forms that are used to express certain functions, then this amounts to a cross-linguistic generalization. I would like to suggest

that every such restriction on the form-function mapping can be considered to be a (generalized) scale. Traditional one-dimensional scales are just a special kind of such cross-linguistic restrictions on form-function mapping. When this limitation—i.e. that a scale has to be one-dimensional—is discarded, and the concept of ordering is replaced by a concept of distance, then the notion of a scale can be nicely generalized to cover many, if not all, restrictions on form-function mapping (cf. Croft 2003: 133-142).

To establish a scale in its generalized conception, it is necessary, first, to establish a cross-linguistic scale of functions; second, to establish a cross-linguistic scale of forms; and, third, to observe a match between the two. Strictly speaking, a cross-linguistic scale, or hierarchy, is the interpretation of any such observed match. These three topics—scales of function, scales of form, and matching them—will be discussed in turn below in Sections 4 to 6. However, first I will introduce a central tool for the generalization of scales, namely the dissimilarity matrix (Section 2), and the concrete example to be used for the discussion of the generalization, namely the “scale of likelihood of spontaneous occurrence” (Section 3).

2. Replacing ordering with distances

A scale in linguistics is normally conceived of as a linear ordering of categories or functions. However, the restriction to a linear structure is neither necessary nor advantageous. Already the perennial issue of whether first person should outrank second on the animacy scale (or vice versa) illustrates that a linear ordering is simply not powerful enough to model linguistic diversity. Probably the only reason for the existence of this focus on linear orders is that such scales are easier to handle and easier to visualize. Further, many scales currently being discussed in the literature only consist of two entities, so that the whole issue of linearity does not arise. However, to generalize the notion of a scale, it seems more fruitful to abandon the principle of a linear scale and open up the possibility for more complex topologies.

One proposal for more complex structures is implicit in the spider-web-like graphs used to display semantic maps (Croft 2003: 133-139; Haspel-math 2003). I have argued elsewhere that such semantic maps can be generalized as dissimilarity matrices (Cysouw 2007; 2010). In a dissimilarity matrix, all pairs of entities in the scale are considered separately, and eval-

uated individually on their similarity. A linear scale is a special case of such a general structure. For example, consider three entities A, B, and C on a linear scale. This translates to similarities by stating that the distance from A to B is the same as the distance from B to C, and both are exactly half the distance from A to C. When the distance from A to C does not exactly match the summed-up distance from A to B to C, the distances do not fit on a linear scale anymore. The principle of the generalization proposed here is that linearity might still exist, but it is not assumed *a priori*. Initially, the pairwise distances are established individually. Only afterwards might it turn out that they reduce to a nicely linear arrangement. However, most of the time they will turn out not to be that easily aligned.

Often—though not in all instances—scales are considered to be inherently directed, i.e. they have a “high end” and a “low end”.² As I will argue in Section 5, this direction is solely caused by the scale of form (cf. Croft 2003: 140-142). Any *direction* of a scale (based on form) is an independent insight from the discovery of the underlying configuration, or *topology*, of the scale (which is based on function). How to practically proceed with the separation of these two issues is the central proposal of this paper.

3. Scale of likelihood of spontaneous occurrence

As an example of the approach proposed here, I will reanalyze data from Haspelmath (1993) on the causative/inchoative alternation. In this paper, Haspelmath addresses the question of how languages mark the predicate in the alternation between an inchoative expression (i.e. an expression without an instigator of the action, like *the water boiled*), and a causative expression (i.e. an expression with an instigator, like *the man boiled the water*). He proposes a scale of the “likelihood of spontaneous occurrence” to explain why—across languages—some predicates tend to be causativized (i.e. the causative is morphologically derived from the inchoative, e.g. German inchoative *enden* vs. causative *beenden*), while others tend to be anticausativized (i.e. the inchoative is derived from the causative, e.g. English inchoative *be destroyed* vs. causative *destroy*). The idea of the scale of likelihood of spontaneous occurrence is that those predicates that are likely to occur spontaneously (i.e. without any human agent) will cross-linguistically tend to be causativized (i.e. the inchoative is the more basic expression), and *vice versa*.

4 *Error! Style not defined.*

Haspelmath investigated the inchoative/causative alternation of 31 meanings in 21 languages. For each of these meanings. The proportion of languages that use a causativization strategy is shown in Table 1.³ The order of the meanings in the table illustrates the idea of the scale of likelihood of spontaneous occurrence, with the least spontaneous meanings at the top (*split, close, break*) and the most spontaneous meaning at the bottom (*dry, freeze, boil, die/kill*).

Table 1. Cross-linguistic proportion of causativizations (adapted from Haspelmath 1993: 104).

<Insert Table 1>

4. Scales of functions

In most current research, scales of categories/functions are either available as hypotheses from earlier research, or established *post-hoc* as the most compelling way to interpret an observed cross-linguistic scale of form (see Section 5) and as such, they can be used as hypotheses for future research. However, scales of function can actually be established independently by using the semantic map approach, though such “scales” will normally not be nicely one-dimensional, needing more effort for their interpretation.

The basic intuition behind the semantic map approach is that cross-linguistic variation in the expression of the functions can be used as a proxy to the relation between the functions themselves. The central assumption made in this approach is that when the expressions of two functions are similar in language after language, then the two functions themselves are similar. Individual languages might (and will) deviate from any general pattern, but when combining many languages, overall the cross-linguistic regularities will overshadow such aberrant cases.

Thus, the similarity between expressions is the basic measure for establishing a semantic map. However, there are two crucially different kinds of similarity between expressions, only one of which will be used to establish the semantic map. For a semantic map, it is important whether two expressions are similar on purely language-specific grounds, i.e. they behave alike according to the grammar of the language. For example, the English verbs *walk* and *enter* behave alike as to the formation of their past forms (*walked, entered*). Likewise, the verbs *buy* and *fight* behave alike in choosing the same kind of past formation (*bought, fought*). This kind of similarity be-

tween expression is purely language-particular and thereby crucially different from cross-linguistic coding strategies. Cross-linguistically, one might say that *walk* and *enter* both use suffixal concatenative morphology to mark past, but *buy* and *fight* use ablaut-like non-linear morphology. Such typological characteristics lead to scales of form to be discussed in the next section (see Cysouw 2010 for a more detailed exposition of this approach using linguistic behavior to establish semantic maps).

Coming back to the inchoative/causative alternation, the English expressions *open* and *close* have some similarity because they use the same inchoative/causative alternation (both verbs do not change their morphology in this alternation, i.e. they use a labile strategy in the terminology of Haspelmath 1993). Likewise, the German expressions *öffnen* and *schließen* are similar because they use the same inchoative/causative alternation, though they use a different strategy from the one found in English (in German both verbs are anticausativizing, as the inchoative form is derived from the causative form by using reflexive morphology: *sich öffnen*, *sich schließen*).

By combining such language-particular similarities from many languages, the similarity between functions/meanings can be approximated. So, in the above example, both English and German use the same (language-particular) construction for “to open” and “to close”, so both languages argue for some similarity between these meanings. However, this is not necessary the case in all the world’s languages. For example, in Hindi the inchoative/causative alternation for the verb “to open” is coded by non-linear morphology (ablaut) *khulnaa/kholnaa*, but the alternation for the verb “to close” it is coded by using the copula-like verbs *honaa* “to be” and *karna* “to do”, viz. *band honaa/band karna*, lit. “be close/do close”. Now, it is possible to approximate the similarity of the meanings “to open” and “to close” by (roughly speaking) taking the average of many such language-particular similarities.

When this procedure is followed for all possible pairs of meanings (see the appendix of Cysouw 2010 for the basic data), this will result in a long list of similarity measures of two meanings. For example, in the case of the meanings investigated by Haspelmath, there are $31 \cdot 30 / 2 = 465$ such pairs of meanings. Such a long list of numbers (a “dissimilarity matrix”) is a generalized scale of meaning. The network in Figure 1 is an attempt to display the structure of the resulting “scale” of meanings. The figure shows a so-called “splits graph” (Bandelt & Dress 1992; Dress & Huson 2004).⁴

Roughly speaking, similar functions will be placed close to each other in the network-like graph. At the upper right side of the figure some of the more spontaneous meanings can be found (e.g. *freeze, dry, burn, boil*, cf. Table 1), and at the opposite site, at the left and lower left, the meanings that typically need an agent are located (e.g. *break, split, open, change*). It is already possible to discern something like the scale of likelihood of spontaneous occurrence. However, this graph also very clearly shows that the spontaneity scale is not the only thing that matters. As one might expect, meaning/function is a highly complex and multidimensional matter (cf. Wälchli & Cysouw 2011), and a multitude of other aspects of meaning are relevant for the similarity of meaning between the meanings investigated.

Another way to depict the structure in the list of pairwise similarities is to use multidimensional scaling.⁵ Shown in Figure 2 are the first two dimensions of a multidimensional scaling for the same data that resulted in the network in Figure 1. Because only the first two dimensions are shown, this display might look easier to interpret, but that is only because much of the complexity of the data is ignored to fit the display into two dimensions. The spontaneity scale can be seen ranging from the upper left to the lower right side in Figure 2 (cf. Table 1). The meanings in the upper left of the figure are highly spontaneous (*boil, freeze, dry*), while the meaning at the lower right typically need an agent (*open, split, close, break, change*).

Figure 1. A scale of function of the 31 meanings in the form of a NeighborNet.

<Insert Figure 1>

Figure 2. First two dimensions of multidimensional scaling of the 31 meanings.

<Insert Figure 2>

5. Scales of form

The constructions that languages use to mark the inchoative-causative alternation are not directly comparable across languages. Take, for example, the German construction using *sich* to mark the inchoative (e.g. *sich öffnen, sich schließen*). This construction is very reminiscent of the Hebrew construction using *hit-* to mark the inchoative (e.g. *hitforer, hit?asef*). Both constructions are of course different in principle—after all, they come from different languages. However, there are various characteristics that make

both constructions alike to some extent. For example, they both explicitly mark the inchoative in relation to the causative. Also, both constructions perform this marking by putting some extra material in front of the lexical verb (though there is of course a difference in morphological status). Further, both the German *sich* construction and the Hebrew *hit-* construction are sometimes considered to be “reflexive” constructions. These three characteristics are cross-linguistically applicable, and in this sense crucially different from the characteristics that are used to establish language-specific construction classes (Haspelmath 2010).

Such cross-linguistically applicable characteristics of expressions are called “strategies” in the typological literature. The tradition of using the term “strategy” in this way probably originated with Keenan & Comrie’s (1977: 64) classic paper on relativization strategies. There are different kinds of strategies, and these different kinds have a rather different status for the comparison of languages, but that topic will not further developed here (see Cysouw 2010). I will here only use so-called “coding” properties that relate to the form in which the language-particular expressions are codified.⁶ In this realm, one can think of characteristics like length of forms, kind of morphological process, or order of elements. The similarity of constructions with respect to such a coding property is here called a *scale of form*.

In Haspelmath’s original 1993 paper on the inchoative/causative alternation, he distinguishes five different coding strategies that languages use to mark the alternation: causative, anticausative, equipollent, labile, and suppletive. Causative constructions are inchoative/causative pairs in which the causative is morphologically overtly derived from the inchoative. Anticausative constructions are the opposite: the inchoative is overtly derived from the causative. Labile constructions are alternations that do not show any overt marking on both inchoative and causative, in contrast to equipollent constructions that have some marking on both. Finally, suppletive constructions are inchoative/causative alternations where there is no (obvious) morphological relation between the two forms. The central opposition in this scale of form is the causative vs. anticausative opposition, the analysis of which led Haspelmath to the spontaneity scale (cf. Table 1).

Many such scales of form can rather easily be approximated by automatically generated measures. Such measures will never be perfect from a linguist’s perspective, but they will get the job done much more quick. For example, consider simply counting the number of Unicode characters used

in the written version of the inchoative and causative forms. A plot for the average (Unicode-based) wordlength of the inchoative vs. the causative is shown in Figure 3. Obviously, these two counts are strongly correlated because in most cases there is regular morphology deriving one from the other, and the counts of characters include the length of the stem. However, there appears to be an interesting cross-linguistic cline in the total length of the meanings. The expressions of “die/kill” tend to be short, while the expression of “develop” and “improve” seem to be long across languages. This cline might be correlated to frequency of use, in that more frequently occurring meanings have shorter expressions. However, quickly checking some online corpora, the lower left to upper right cline in Figure 3 does not seem to correlate well with pure token frequency.

More relevant to the current topic is the upper left to lower right cline in Figure 3, which represents the different in length between the causative and the inchoative form. Meanings in which, across languages, the causative is longer than the inchoative should correspond to those meaning that have a preference for causativization, i.e. they should be high on the spontaneity scale shown in Table 1. Unsurprisingly, as shown in Figure 4 the spontaneity scale strongly correlates with the average causative-minus-inchoative character count ($r = 0.89$). Actually, the only meaning clearly being off on this correlation is “die/kill”, which is probably an effect of the imprecise estimate on the spontaneity scale (cf. note 3), and not so much an error of the approximation of counting characters. It does seem to make more sense to place “die/kill” somewhere on the higher middle of the spontaneity scale (as suggested by the counts of characters) than to place it completely on top (as suggested by the spontaneity scale in Table 1). After all, dying is indeed commonly a spontaneous activity, though it is not that uncommon to be induced by an agent.

Figure 3. Average number of characters used to mark inchoative (x-axis) and causative (y-axis).

<Insert Figure 3>

Figure 4. Spontaneity scale approximated by average length difference between the inchoative and causative forms.

<Insert Figure 4>

6. Matching form and function

Given a scale of function and a scale of form, the basic idea now is to investigate how these two scales match. In general, it is not immediately obvious how this should be done, but for specific cases there are many nice techniques to visualize such correlations and investigate their statistical significance. The examples discussed below illustrate some of the possibilities.

In Figure 5, the multidimensional scaling from Figure 2 is used as the basis to display the scale of function (i.e. the semantic map). The scale of form (i.e. the length difference of the actual forms) is shown as an overlay over this display. This overlay is like a geographic map using contour lines (technically called “isohypses”) to indicate elevation. The level of elevation is defined by the total difference of characters between the inchoative and causative forms throughout all 21 languages.⁷ The scale of spontaneity can be clearly seen ranging from the highest point for “boil” and “freeze” at +25 to the lowest point for “close” at –25. This indicates that the scale of function (cf. Section 4, approximated by the location of the points in Figure 5) already includes the scale of spontaneity to some extent. This is noteworthy, because the scale of functions was made without any knowledge about the length of the forms, nor about the causativization patterns.

Figure 5. Summed difference in length between causative and inchoative forms throughout all languages, shown as an overlay over the scale of function.

<Insert Figure 5>

It is also possible to give a more precise analysis of how strong the overlap between the two scales is. Statistically, the question here is to which extent the scale of function can be explained by the scale of form. This problem is somewhat alike to a multivariate analysis of variance, if it were not for the fact that the variable to be explained (the scale of function) is of a rather unusual kind, namely a dissimilarity matrix. Recent work in bioinformatics (Zapala & Schork 2006) fortunately presents a solution for this particular problem.⁸ As shown in Table 2, the length of the causative and the length of the inchoative explain about 36% of the distances in the scale of function. Or, more to the point, the difference in length between inchoative and causative explains about 21% of the variation, and the sum 15% (as shown in Table 3).

Table 2. ANOVA of distance matrix by length of forms.

<Insert Table 2>

Table 3. ANOVA of distance matrix by length difference and length sum.

<Insert Table 3>

A different approach to the correlation between the scale of function and the scale of form is by using matrix correlation. Basically, the idea is to also consider the scale of form as a dissimilarity matrix and then correlate the form matrix with the function matrix. To reformulate the measurements of form (in the current case, these measurements are the average length of the inchoative and the causative expressions) into a dissimilarity matrix, all pairs of measurements have to be compared individually. As a dissimilarity, one can, for example, simply take the Euclidean distance between the measurements for each pair. This dissimilarity in effect represents the linear distance between the words as shown in the configuration of Figure 3. The length of a direct line between two words in that figure is the same as the Euclidean distance. Mathematically defined, this amounts to taking the dissimilarity as defined in (1).

$$(1) \quad d(A,B)=\sqrt{(\text{Inch}_A-\text{Inch}_B)^2+(\text{Caus}_A-\text{Caus}_B)^2}$$

For example, “boil” has an average length of 5.52 characters for the inchoative and 7.05 for the causative. Likewise, “freeze” has an average length of 6.62 for the inchoative and 8.29 for the causative. Taking the Euclidean distance between the point (5.52, 7.05) and (6.62, 8.29) results in a dissimilarity between “boil” and “freeze” of 1.66. Doing these calculations for all pairs results in a dissimilarity matrix of form. Figure 6 shows the correlation between this scale of form and the scale of function. Each point in this figure represents one pair of meanings, plotting the dissimilarity of function against the dissimilarity of form. The figure already shows a rather nice correlation, which can also be shown to be statistically significant ($r = 0.40$, Mantel test $p < 0.0001$).⁹

Figure 6. Correlating form and function dissimilarities.

<Insert Figure 6>

7. Conclusions

The following points summarize the proposals put forward in this paper:

- A linguistic scale (or hierarchy) consists of three parts: a scale of functions, a scale of form, and a match between the two.
- Both the scales of function and the scales of form are not necessarily linear. They can be internally structured in complex ways. Yet such complex structures are inherently interesting and still represent strong restrictions on the probabilities of linguistic variation.
- The most general description of the internal structure of these scales takes the form of dissimilarity matrices, which might boil down—under special circumstances—to a linear structure.
- The match between form and function is a kind of matrix correlation, though other methods might also be used. However, this is an area where much work has to be done to elucidate which approaches are most suitable for linguistic typology.

Acknowledgements

This paper was written while I was at the Max Planck Institute for Evolutionary Anthropology in Leipzig. A revised version was prepared while I was funded by the ERC starting grant “QuantHistLing”. I thank Martin Haspelmath for in-depth discussion of the paper before the preparation of the revision.

Notes

1. The term “scale” is used here synonymously to what is also known as an “implicational hierarchy”, “markedness hierarchy” or simply “hierarchy” in linguistics.
2. The combination of a linear scale with an inherent direction is mathematically equivalent to a so-called “total order”. An easy way to remedy the problem of having parts of the scale that are not (clearly) ordered relative to each other (like first and second person in the animacy hierarchy) is to allow for some parallelism in the ordering (thus deviating from strict linearity). Such a model is mathematically speaking an example of a so-called “partial order”.
3. The proportion of causatives reported in Table 1 is calculated by dividing the number of languages that causativize the predicate (C) by the number of the languages that either causativize or anticausativize it (C+A), ignoring those

languages that use different strategies (Haspelmath further distinguishes suppletive, labile and equipollent alternations, which will not be used here). This method of calculation is different from the proportions reported on by Haspelmath (who lists the fraction A/C). Further, this fraction is noteworthy in the case of die/kill, as most languages use a suppletive strategy for this meaning, making the proportion reported here (1.00) somewhat superficial (because A=0 and C=3, so $C/(C+A)=1$).

4. The particular splits graph shown in Figure 1 is a NeighborNet made by the program *SplitsTree* (Huson and Bryant 2006). See Bryant et al. (2005) for an introduction to this approach with some examples from linguistics.
5. For all multidimensional scaling in this paper I used the function *cmdscale* from the statistical environment R (R Development Core Team, 2007).
6. Besides coding properties, Keenan (1976) also distinguishes behavioral properties of expressions in complex constructions as another kind of properties. Note that the German/Hebrew example discussed above included yet another kind of cross-linguistic strategy. The impression that both the German *sich* construction and the Hebrew *hit-* construction are “reflexive” constructions can be formalized by including reference to a “prototypical” element in the realm of meaning. For example, constructions from different languages are both reflexive-like when they both at least code for the meaning “rise”.
7. To make such a map, it is first necessary to make an interpolation over the measurements of elevation at the points as defined by the multidimensional scaling. It is not trivial to make such an interpolation, because the points are rather unequally distributed. To make an interpolation, I used a geostatistical technique called “kriging” as implemented by the function *krige.conv* in the R package *geoR* (Ribeiro Jr and Diggle 2001), with the parameter settings *s2=1* and *phi=10*. On this basis, the isohypses were drawn using the *contour* function.
8. The multivariate ANOVAs shown in Table 2 and Table 3 were calculated by using the function *adonis* in the R package *vegan* (Oksanen et al. 2007).
9. The Mantel test (Mantel 1967) was performed using the function *mantel.test* in the R package *APE* (Paradis et al. 2004).

References

- Bandelt, H.J., and A.W. Dress
 1992 Split decomposition: a new and useful approach to phylogenetic analysis of distance data. *Molecular Phylogenetics and Evolution* 1(3): 242-252.

- Bickel, Balthasar, and Alena Witzlack-Makarevich
this volume Referential scales and case alignment: reviewing the typological evidence.
- Bryant, David, Flavia Filimon, and Russell D. Gray
2005 Untangling our past: Languages, trees, splits and networks. *The Evolution of Cultural Diversity: A Phylogenetic Approach*, Ruth Mace, Clare J. Holden, and Stephan Shennan (eds.), 67-84. London: UCL.
- Croft, William
2003 *Typology and Universals*. (Cambridge Textbooks in Linguistics). Cambridge: Cambridge University Press.
- Cysouw, Michael
2007 Building semantic maps: the case of person marking. *New Challenges in Typology*, Bernhard Wälchli and Matti Miestamo (eds.), ?? Berlin: Mouton de Gruyter.
- Cysouw, Michael
2010. Semantic maps as metrics on meaning. *Linguistic Discovery*. 8(1): 70-95.
- Dress, Andreas W. M., and Daniel H. Huson.
2004. Constructing Splits Graphs. *IEEE Transactions on Computational Biology And Bioinformatics* 1(3): 109 - 115.
- Haspelmath, Martin
1993 More on the typology of inchoative/causative verb alternations. *Causatives and Transitivity*, Bernard Comrie and Maria Polinsky (eds.), 87-120. Amsterdam: Benjamins.
- Haspelmath, Martin
2003 The geometry of grammatical meaning: Semantic maps and cross-linguistic comparison. In *The New Psychology of Language: Cognitive and Functional Approaches to Language Structure*, Michael Tomasello (ed.), Volume 2, 211-242. Mahwah, NJ: Erlbaum.
- Haspelmath, Martin
2010 Comparative concepts and descriptive categories in cross-linguistic studies. *Language* 86(3). 663-687.
- Huson, Daniel H., and David Bryant
2006 Application of phylogenetic networks in evolutionary studies. *Molecular Biology and Evolution* 23(2): 254-267.
- Keenan, Edward L.
1976 Towards a universal definition of 'subject'. In *Subject and Topic*, Charles N. Li (ed.), 303-333. New York, NY: Academic Press.
- Keenan, Edward L., and Bernard Comrie
1977 Noun phrase accessibility and universal grammar. *Linguistic Inquiry* 8(1): 63-99.

Mantel, Nathan

- 1967 The detection of disease clustering and a generalized regression approach. *Cancer Research* 27(2): 209-220.

Oksanen, Jari, Roel Kindt, Pierre Legendre, Bob O'Hara, and M. Henry H. Stevens

- 2007 vegan: Community Ecology Package. *R package*.

Paradis, E., J. Claude, and K. Strimmer

- 2004 APE: analyses of phylogenetics and evolution in R language. *Bioinformatics* 20, 289-290.

R Development Core Team

- 2007 *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Ribeiro Jr, Paulo Justiniano, and Peter J. Diggle

- 2001 geoR: A package for geostatistical analysis. *R News* 1, no. 2: 15-18.

Silverstein, Michael

- 1976 Hierarchy of features and ergativity. *Grammatical Categories in Australian Languages*, R M W Dixon (ed.), 112-171. Australian Institute of Aboriginal Studies: Canberra.

Wälchli, Bernhard, and Michael Cysouw

- 2011 Lexical typology through similarity semantics: Toward a semantic map of motion verbs. *Linguistics* (forthcoming).

Zapala, MA, and NJ Schork

- 2006 Multivariate regression analysis of distance matrices for testing associations between gene expression patterns and related variables. *Proceedings of the National Academy of Sciences of the United States of America* 103(51): 19430-19435.