Spatial terms reflect near-optimal spatial categories

Naveen Khetarpal (khetarpal@uchicago.edu)[*]
Asifa Majid (asifa.majid@mpi.nl)[†]
Terry Regier (regier@uchicago.edu)[*]

[*]Department of Psychology, University of Chicago, Chicago, IL 60637 USA
[†]Max-Planck Institute for Psycholinguistics, 6500 AH Nijmegen, The Netherlands

Abstract

Spatial terms in the world's languages appear to reflect both universal conceptual tendencies and linguistic convention. A similarly mixed picture in the case of color naming has been accounted for in terms of near-optimal partitions of color space. Here, we demonstrate that this account generalizes to spatial terms. We show that the spatial terms of 9 diverse languages near-optimally partition a similarity space of spatial meanings, just as color terms near-optimally partition color space. This account accommodates both universal tendencies and cross-language differences in spatial category extension, and identifies general structuring principles that appear to operate across different semantic domains.

Keywords: Spatial terms; semantic universals; linguistic relativity; language and thought; cognitive modeling.

Spatial terms across languages

The categories picked out by spatial terms differ across languages, but also reveal universal tendencies (Bowerman & Pederson, 1992; Bowerman, 1996; Talmy, 2000; Bowerman & Choi, 2001). Melissa Bowerman designed a set of 71 spatial scenes (see a sample of 10 in Figure 1), which have been used to elicit spatial terms in a variety of languages. Levinson et al. (2003) analyzed the spatial terms applied to these scenes by speakers of 9 unrelated languages. This study and others confirm the general picture of universal tendencies co-existing with language-specific categorization patterns.
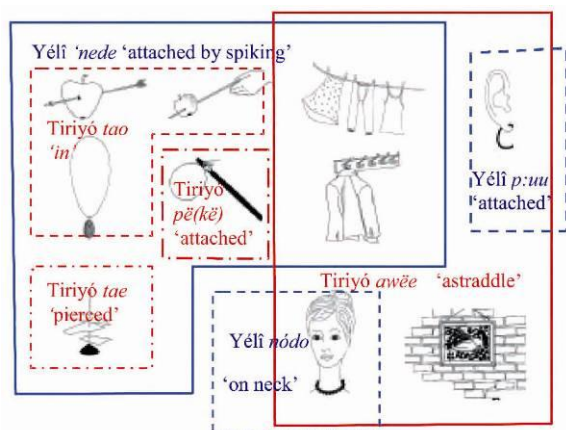

Figure 1: 10 spatial scenes, as categorized in 2 languages: Tiriyó and Yélî-Dnye. Source: Levinson et al. (2003).

The analogy with color

As Levinson et al. (2003) noted, a similar situation obtains with respect to color naming across languages: there are both universal tendencies and cross-language differences. The upper panel of Figure 2 shows a grid of color stimuli often used in color naming research, while the lower panel shows the color naming system of a particular language, Wobé, mapped against that grid: each colored region is the extension of a color term. There is evidence consistent with the universalist view that color categories across languages are organized around a set of universal *focal colors* (Berlin & Kay, 1969; Regier et al., 2005). Yet there is also evidence consistent with the opposing relativist view that category extension is determined by language, not just by these focal colors (Roberson et al., 2005).
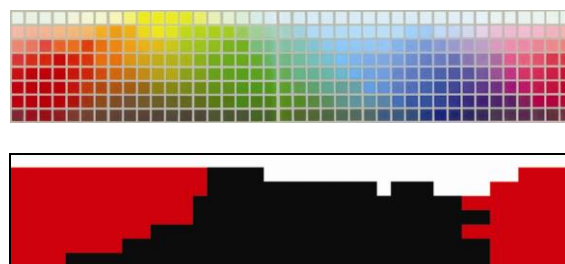

Figure 2: Color naming grid (upper panel), and the color naming system of Wobé, spoken in Ivory Coast, mapped against that grid (lower panel).

A proposal that may resolve this tension was advanced by Jameson and D'Andrade (1997). They pointed out that perceptual color space is irregularly shaped. In the grid of Figure 2, hue appears to vary evenly along the horizontal axis, and lightness evenly along the vertical axis, but this obscures the third dimension of color: saturation. In this grid, each color is shown at the maximum possible saturation for that hue/lightness pair – and that maximum possible saturation varies unevenly across the grid, meaning that the "outer skin" of color space is unevenly shaped. Jameson and D'Andrade suggested that color names across languages partition this uneven space such that names are maximally informative about color. Regier et al. (2007) formalized this proposal and tested it against data from the World Color Survey (WCS), a color naming database drawn from 110 languages of non-industrialized societies worldwide. They defined the *well-formedness* of a given

categorical partition of color space to be the extent to which such a partition maximizes similarity within categories, and minimizes it across categories (Garner, 1974). Concretely, they took the similarity between two colors i and j, $sim(i,j)$, to be a Gaussian function of the distance between colors i and j in a standard 3-dimensional perceptual color space, CIELAB. Given this, they defined:

$$S_w = \sum_{\substack{(i,j): \\ cat(i)=cat(j)}} sim(i,j)$$

$$D_a = \sum_{\substack{(i,j): \\ cat(i)\neq cat(j)}} (1 - sim(i,j))$$

$$W = S_w + D_a$$

Here S_w is an overall measure of within-category similarity, obtained over all unique pairs of colors i,j that fall in the same category; D_a is an overall measure of across-category dissimilarity, over all unique pairs that fall in different categories; and W is well-formedness. Regier et al. used steepest ascent in well-formedness to create theoretically optimal color naming systems with n=3,4,5,6 categories, and found that these theoretical optima were similar to actual color naming systems found in some languages. They also found that most color naming systems in the WCS were *near-optimal*, in the sense that they were more well-formed than a natural comparison class of unattested systems of comparable complexity. The top panel of Figure 3 shows the color naming system of Berinmo (Roberson et al., 2000), while the middle and bottom panels show hypothetical systems obtained by rotating the Berinmo system by different amounts in the hue dimension. If color naming in Berinmo is near-optimal, we would expect the attested system to show higher well-formedness than any hypothetical system derived from it by rotation to any position along the hue dimension. This prediction was confirmed.
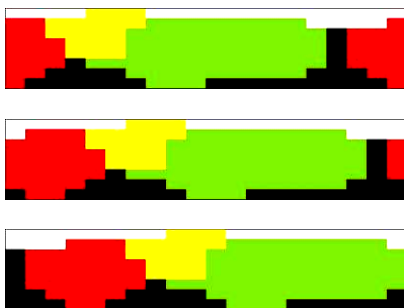

Figure 3: The Berinmo color naming system unrotated (top), and rotated four (middle) and eight (bottom) hue columns. Adapted from Regier et al. (2007).

The same rotation analysis was applied to each of the 110 languages of the WCS. Figure 4 shows well-formedness (linearly transformed to the range 0-1) as a function of the number of columns of rotation in the hue dimension,

averaged across all 110 languages of the WCS. It can be seen that there is a clear tendency across languages for the attested, unrotated system to "fit the shape" of color space: to be more well-formed than unattested variants of comparable complexity, derived by rotation in hue.[1]
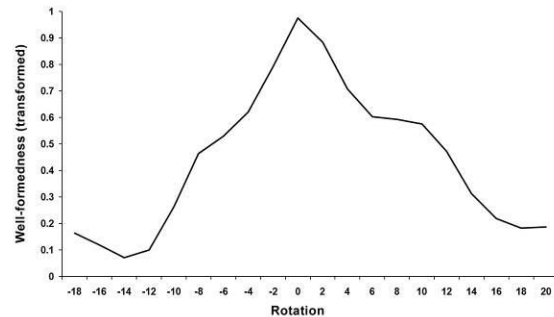

Figure 4: Well-formedness averaged across all 110 WCS languages, as a function of rotation. Adapted from Regier et al. (2007).

In this sense, color naming across languages tends to be near-optimal, a finding that accounts for universal tendencies in color naming. However the same account also suggests how language may play a role in determining category boundaries. There are often many highly well-formed systems, sometimes varying subtly from one another, and linguistic convention may select from among this set of high-ranked candidate systems. This would account for the mixed picture of universal tendencies and some language-specificity in the demarcation of boundaries. In this manner, the well-formedness account suggests a specific middle ground between nature and nurture in color naming.

Are named spatial categories near-optimal?

Can these ideas also account for spatial language, in which we see a similarly mixed pattern of universals and cross-language variation? The present study tests this question.

In place of the color naming grid and the naming data of the World Color Survey relative to that grid, we used the 71 line-drawn spatial scenes of the TOPOLOGICAL RELATIONS PICTURE SERIES designed by Melissa Bowerman (see e.g. Bowerman & Pederson, 1992) and illustrated in part in Figure 1, as named by speakers of 9 unrelated languages: Basque, Dutch, Ewe, Lao, Lavukaleve, Tiriyó, Trumai, Yélî-Dnye, and Yukatek (Levinson et al., 2003). In place of similarities based on the CIELAB color space, we obtained pairwise similarities among these spatial scenes based on pile-sorting of these scenes by speakers of English and Dutch.[2] Finally, in place of rotation in the hue dimension, we used rotation in a similarity space derived from the pile-

---

[1] For 82 of the 110 languages of the WCS, the attested system yielded higher well-formedness than any rotation of that system.

[2] These two languages were chosen for convenience only. Future research can usefully explore the sensitivity of our results to the native language of pile-sorters.

sorting. We wished to determine whether the spatial naming systems of these 9 languages are near-optimal in a manner analogous to that demonstrated for color naming systems.

Methods

We describe here the methods undertaken to perform this test: the treatment of the naming data, the measurement of similarities, the well-formedness function, and the rotation method.

Naming data

The naming data had been collected previously (see Levinson et al., 2003). For each of the 9 languages, we recorded, for each of the 71 spatial scenes, the modal spatial adposition for that scene — i.e. the spatial term that was used by the largest number of speakers of the language to name that scene. Ties were broken by random choice. The resulting array of names was taken to be that language's spatial naming system for those scenes.

Similarity judgments

It would have been prohibitively time-consuming to collect pairwise similarity judgments for all possible pairs of the 71 scenes. Therefore we instead asked participants to sort the scenes into piles (cf. Malt, Sloman, Gennari, Shi, & Wang, 1999; Rosenberg & Kim, 1975). The pile-sorting data were then used to derive a measure of similarity between each pair of spatial relations.

Participants. 24 native speakers of Dutch and 22 native speakers of English participated in the sorting task, at the MPI in Nijmegen and the University of Chicago, respectively.

Materials. The 71 line drawings were printed individually onto 6.5×6.5 cm cards. The FIGURE objects were drawn in orange and the GROUND objects were drawn in black. (The FIGURE is the object which is to be thought of as located relative to the GROUND object.) The tabletop on which participants sorted the cards into piles was $115(W) \times 60(D)$ cm.

Procedure. Participants were verbally instructed to "group the cards into piles so that in each pile, the spatial relation of the orange object to the background black object is similar." The directionality of this relation was emphasized and reiterated. To help ensure that they would judge similarity with respect to spatial arrangement and not with respect to the identity of the objects involved, participants were further told that they "may see the same objects on multiple cards" and instructed to "not put cards together just because they contain an image of the same object or objects." Participants were instructed in their native language; the English instructions from which these excerpts come were translated to Dutch, and back-translated to English in an effort to minimize possible differences in meaning.

Participants were not restricted in the amount of time they had to complete the task. They were also told that they were free to make as many piles as they needed, and that there was no restriction on the number of items allowed in a single pile.

Analysis. The similarity between each pair of spatial scenes was determined by the proportion of all participants who placed those two cards in the same pile. For example, if all 46 participants placed card i and card j in the same pile, then the similarity between i and j would be $sim(i,j) = 46 \div 46 = 1$; if 23 of the participants placed these two cards in the same pile, then the similarity between i and j would be $sim(i,j) = 23 \div 46 = 0.5$, and so on.

Well-formedness

The well-formedness function W was the same as that of Regier et al. (2007), as described above, with the adjustment that the function $sim(i,j)$ for any pair of spatial scenes i and j was determined by pile-sorting as just described.

Rotation

We wished to determine whether the linguistic spatial systems represented by our naming data were near-optimal in the same sense as color naming systems – that is, whether these attested systems were more well-formed than a natural comparison class of unattested systems of comparable complexity, derived by rotation. Rotation is an appropriate means of generating hypothetical systems comparable to the original, since the number of categories, the number of stimuli per category, and (to the extent allowed by the irregularity of the space) the location of the categories relative to each other are all maintained. Performing rotation-based comparisons requires a similarity space within which rotation can occur. Therefore, we (1) created a similarity space from the pile-sort similarity data, (2) rotated the naming systems in that space, and (3) compared well-formedness of attested and rotated systems. We describe these steps in turn.

Creating a similarity space. The pairwise similarities obtained by pile-sorting were provided as input to a multidimensional scaling algorithm, ALSCAL, which produced a similarity space from them (3-dimensions, stress = 0.180; only minimal further reduction in stress was obtained with more dimensions). Each of the 71 spatial scenes is represented as a point in this space, and the distances between points approximate the dissimilarity between the corresponding spatial scenes. Figure 5 shows the 71 points plotted in 2 dimensions of the resulting 3-dimensional space. The points are irregularly positioned in this space, by rough analogy with the irregularity of color space that appears to drive patterns of color naming across languages (Jameson & D'Andrade, 1997; Regier et al., 2007).
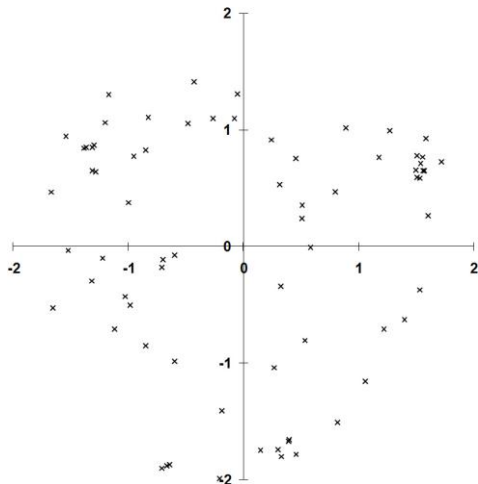
Figure 5: The 71 spatial scenes plotted as points in 2 dimensions of a 3-dimensional similarity space.

Rotating in similarity space. In the color domain, rotating was an apparently simple matter of shifting a color naming system "rightward" or "leftward" in the color naming grid, along the hue dimension. In the spatial domain, rotating naming systems is less straightforward, since there exists no explicit grid – and when the points are plotted in the MDS-produced similarity space, as in Figure 5, the points are not arranged in an implicit grid either. Moreover, in the similarity space there are 3 dimensions around which one could rotate.

Therefore, we rotated a given spatial naming system in this similarity space as follows. We first selected one of the 3 dimensions of the space as the axis around which we would rotate, and then rotated all points in the 3-D space around that axis by the same amount. Specifically, call the set of all 71 points P. We rotated the set of points P around the selected axis by some angle α, resulting in a new set of points P'. In general, these points P' did not coincide with the original points P – it is in this sense that there is no implicit grid. Nonetheless we wished to relabel the original points P, which represent the spatial scenes in our stimulus set, according to the labels on nearby rotated points in P'. We therefore identified a bijection (one-to-one correspondence) between points in P and P' that near-minimized the total distance between corresponding points in the two sets[3], and assigned to each point in P the label of the corresponding point in P'. We took the resulting new labeling of the points P to be the rotated labeling at angle α, around the selected dimension. We rotated each spatial naming system by increments of 15° all the way around the

---

[3] Specifically, we considered the points P in random order, and assigned each point $p \in P$ to the nearest as-yet-unassigned point $p' \in P'$. We repeated this process 10^6 times with different random orderings of P, and chose the mapping that yielded the minimum overall distance between corresponding points in P and P'. We also explored a variant of this procedure in which we instead chose the mapping that maximized well-formedness. The overall results were qualitatively the same.

circle (i.e. 0°, 15°, 30°, 45°, up through 345°), and did so separately around each of the three dimensions of the similarity space.

This procedure yields, for each dimension of the similarity space, 23 hypothetical naming systems that are comparable to the original naming system in that they have the same number of categories, the same number of stimuli per category, and roughly the same arrangement of the categories relative to each other, as in the original. Some distortion of the shape of the categories is introduced by the rotation process, because of the necessity of coercing points in P' to nearby points in P – but note that something similar happens in the case of color rotation as well: when a color naming system is shifted "rightwards" or "leftwards" in the grid, the labels are shifted to other colors that differ not only in hue, but also in saturation because of the irregular shape of the space, resulting in some distortion of overall category shape. Critically, in both color and space, if there were no such distortion of the categories under rotation, well-formedness would not differ under rotation – thus, these distortions are what allow us to determine how well a given labeling system "fits" a given similarity space.

Comparing well-formedness. For each of the 9 languages in our sample, we rotated that language's spatial naming system around each of the 3 dimensions of the similarity space, as described above. We then determined the well-formedness of the original system and of each rotated variant.

Results

Figure 6 displays the well-formedness of the Lao spatial naming system, and of hypothetical variants derived from it by rotation around each of the 3 dimensions (arbitrarily labeled x, y, and z) of the similarity space. The attested system is shown at 0° rotation. The attested system exhibits higher well-formedness than any rotation around the x and z axes, and higher well-formedness than any rotation around the y axis with one exception, which is immediately adjacent to the attested system.
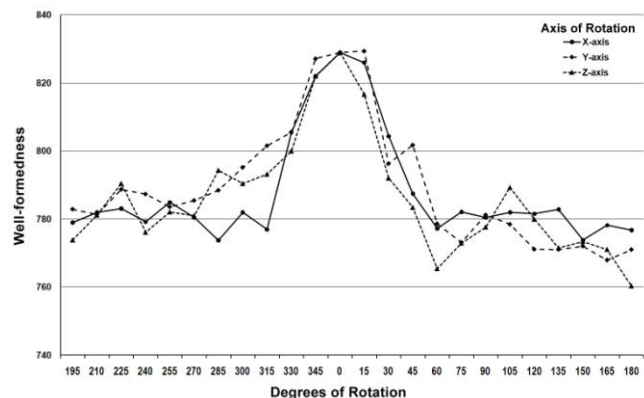

Figure 6: Well-formedness of Lao when rotated about each of the 3 axes.

These results are qualitatively similar to those obtained by Regier et al. (2007) for color naming, and they suggest that in Lao at least, spatial terms are near-optimal in the sense that the categorical partition of similarity space that they define is more well-formed than is (nearly) any member of a natural comparison class of hypothetical partitions of comparable complexity. Thus, the Lao spatial naming system appears to "fit" the shape of the similarity space, much as color naming systems appear to fit the shape of color space.
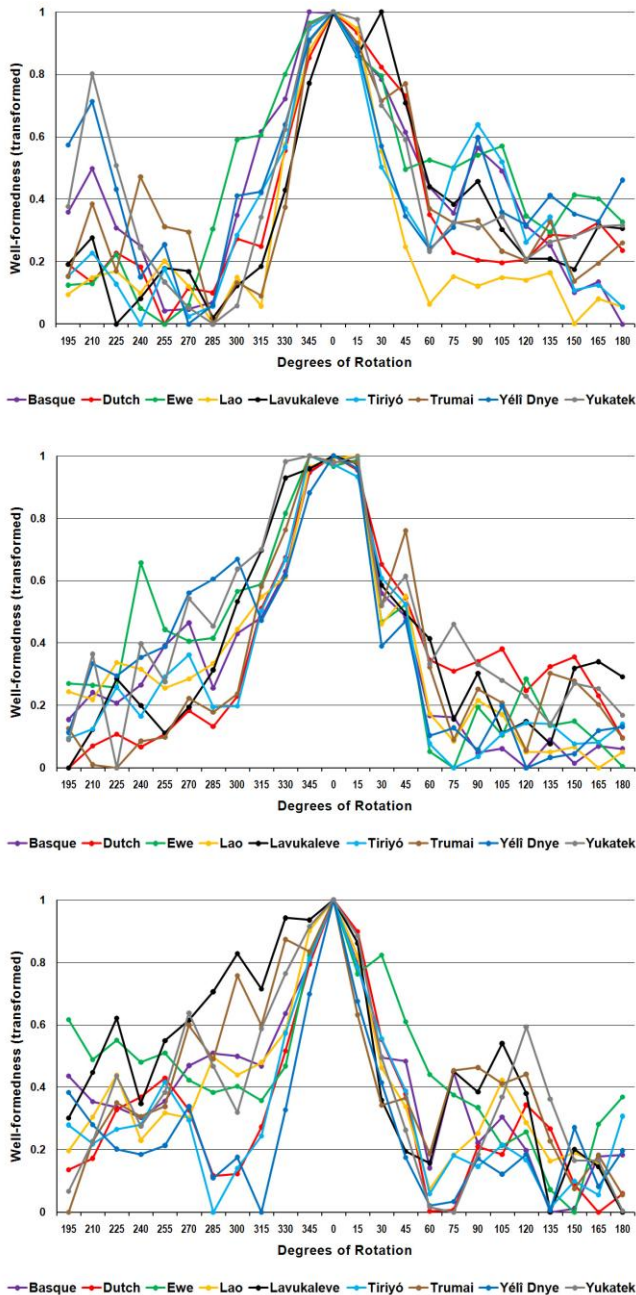

Figure 7: Well-formedness of all 9 languages when rotated about the *x* (top panel), *y* (middle panel), and *z* (bottom panel) axes.

Figure 7 shows that this pattern holds for all 9 languages in our dataset: the top, middle, and bottom panels in the figure show the well-formedness of each of the 9 languages when rotated about the *x*, *y*, and *z* axes, respectively. The attested systems have higher well-formedness than almost all hypothetical variants of them that were considered,[4] and are in this sense near-optimal – which helps to explain universal tendencies in spatial naming. At the same time, these results suggest a possible role for local linguistic convention in determining spatial categories. The fact that the spatial systems of these languages are quite different, yet are all near-optimal, suggests that there may be many highly well-formed partitions of the similarity space, and that linguistic convention may select from among these highly-ranked systems.

Conclusions

There are universal tendencies in the categories picked out by spatial terms across languages, and there is also substantial cross-language variation. The universal tendencies suggest that spatial terms reflect universal cognitive structure of some sort, while the cross-language variation suggests support for the opposing relativist view that linguistic convention plays a central role in determining the extension of the categories found in any given language. In this language-and-thought debate, as in many others, the evidence offers partial support for both positions, and full support for neither.

We have argued that this empirically complex picture can be accounted for in a theoretically straightforward way. Levinson et al. (2003) suggested an analogy between spatial terms and color terms, and we have pursued this idea here. Specifically, we have shown that just as color terms near-optimally partition color space (Regier et al., 2007), so spatial terms near-optimally partition an underlying similarity space. On this view, the universal tendencies in spatial naming result from irregularities in the similarity space, and the cross-language variation reflects the fact that there are often several nearly equally good ways to carve up this space.

There is a potential criticism of this argument that we anticipate. It might be claimed that the analogy with color is flawed, since the irregularity in color space reflects something fundamental, whereas the irregularity in the spatial similarity space may not. Specifically, in color, the irregularity of the space reflects the fact that the maximum obtainable saturation varies unevenly across hue/lightness combinations. In contrast, it may be claimed, the irregularity of our spatial similarity space may primarily reflect the particular sampling of the space of all possible spatial relations that the 71 scenes happen to represent. The grid in the case of color ensures that the sample of colors

---

[4] Two of the nine languages – Dutch and Yélî-Dnye – have peaks at exactly 0˚ rotation around all 3 axes; results for other languages are comparable to those for Lao. Further research is needed to determine whether the Dutch results are unusually clean because the pile-sorting was performed in part by Dutch-speakers.

considered is a systematic one; we have no such assurance in the case of the spatial stimuli we have used here.

It is true that the sample of stimuli we have relied on here appears to be less systematically arrayed than in the case of color – this may be unavoidable in a semantic domain as complex as topological spatial relations, which can capture not just the actual location of an object, but also its method of attachment to another object, whether it is a piece of clothing, etc. And it is likely that some of the irregularity in our similarity space indeed reflects the manner in which the full space of spatial relations was sampled rather than any sort of universal constraint on spatial cognition. Another obvious limitation is that our similarity space is based on responses from speakers of only two closely related languages, English and Dutch. However, despite this, we suspect that some of the structure of the resulting space does indeed reflect universal constraints on the way humans think about space. Tellingly, the spatial systems of 9 diverse languages (one was Dutch, but the rest were unrelated to either English or Dutch) appear near-optimal when assessed in this space. It is clear why we would obtain these results if the structure of the space reflects at least in part universal aspects of human spatial cognition – and it is not clear why we would obtain them otherwise.

Many questions are left open by the results we have reported here. How important is it to have a systematic sample of spatial stimuli? Would we obtain different results if we tested for near-optimality using a method other than rotation? Just how universal is the similarity space: are there detectable differences in the pile-sorting of English and Dutch speakers that are traceable to their respective systems of spatial terms? Why do languages have different numbers of spatial terms, partitioning semantic space more or less finely? Perhaps most interestingly, when we see that there are hypothetical systems that are slightly more well-formed than the attested system from which they were derived, are there actual spatial naming systems that look like those high-ranked hypothetical systems? One might expect this on the account we have presented here. Some of these questions can be answered through further analysis of existing data, while others will require more data. Yet whatever the eventual answers to these questions, we hope the results we have presented help to clarify the interplay of universal and language-specific forces in the naming of spatial relations across languages.

Acknowledgments

References

Berlin, B. & Kay, P. (1969). *Basic color terms: Their universality and evolution*. Berkeley: University of California Press.

Bowerman, M. (1996). Learning how to structure space for language: A cross-linguistic perspective. In P. Bloom, M. Peterson, M. Garrett, & L. Nadel (Eds.) *Language and space* (pp. 385–436). Cambridge, MA: MIT Press.

Bowerman, M. & Choi, S. (2001). Shaping meanings for language: Universal and language-specific in the acquisition of spatial semantic categories. In M. Bowerman and S. Levinson (Eds.) *Language acquisition and conceptual development* (pp. 475–511). Cambridge: Cambridge University Press.

Bowerman, M. & Pederson, E. (1992). Cross-linguistic studies of spatial semantic organization. In *Annual Report of the Max Planck Institute for Psycholinguistics* 1992 (pp. 53-56).

Garner, W. R. (1974). *The processing of information and structure*. Potomac, MD: L. Erlbaum Associates

Jameson, K. & D'Andrade, R. (1997). It's not really red, green, yellow, blue: An inquiry into cognitive color space. In C. L. Hardin & L. Maffi (Eds.), *Color categories in thought and language* (pp. 295–319). Cambridge, UK: Cambridge University Press.

Levinson, S. C. & Meira, S. (2003). Natural concepts in the spatial topological domain—adpositional meanings in crosslinguistic perspective: An exercise in semantic typology. *Language, 79,* 485-516.

Malt, B. C., Sloman, S. A., Gennari, S., Shi, M., & Wang, Y. (1999). Knowing versus naming: Similarity and the linguistic categorization of artifacts. *Journal of Memory and Language, 40,* 230-262.

Regier, T., Kay, P., & Cook, R.S. (2005). Focal colors are universal after all. *PNAS, 102,* 8386-8391.

Regier, T., Kay, P., & Khetarpal, N. (2007). Color naming reflects optimal partitions of color space. *PNAS, 104*, 1436-1441.

Roberson, D., Davies I. & Davidoff, J. (2000). Color categories are not universal: Replications and new evidence from a stone-age culture. *Journal of Experimental Psychology: General*, 129, 369-398.

Roberson, D., Davidoff, J., Davies, I.R.L., & Shapiro, L.R. (2005). Color categories: Evidence for the cultural relativity hypothesis. *Cognitive Psychology, 50*, 378-411.

Rosenberg, S. & Kim, M. P. (1975). The method of sorting as a data-gathering procedure in multivariate research. *Multivariate Behavioral Research, 10,* 489-502.

Talmy, L. (2000). How language structures space. In L. Talmy (Ed.) *Toward a cognitive semantics, Volume 1* (pp. 177-254). Cambridge, MA: MIT Press.