



Cognitive Science (2012) 1–14
Copyright © 2012 Cognitive Science Society, Inc. All rights reserved.
ISSN: 0364-0213 print / 1551-6709 online
DOI: 10.1111/cogs.12000

Representing Spatial Structure Through Maps and Language: Lord of the Rings Encodes the Spatial Structure of Middle Earth

Max M. Louwerse, Nick Benesh

Department of Psychology, University of Memphis

Received 15 October 2010; received in revised form 2 October 2011; accepted 30 April 2012

Abstract

Spatial mental representations can be derived from linguistic and non-linguistic sources of information. This study tested whether these representations could be formed from statistical linguistic frequencies of city names, and to what extent participants differed in their performance when they estimated spatial locations from language or maps. In a computational linguistic study, we demonstrated that co-occurrences of cities in Tolkien's *Lord of the Rings* trilogy and *The Hobbit* predicted the authentic longitude and latitude of those cities in Middle Earth. In a human study, we showed that human spatial estimates of the location of cities were very similar regardless of whether participants read Tolkien's texts or memorized a map of Middle Earth. However, text-based location estimates obtained from statistical linguistic frequencies better predicted the human text-based estimates than the human map-based estimates. These findings suggest that language encodes spatial structure of cities, and that human cognitive map representations can come from implicit statistical linguistic patterns, from explicit non-linguistic perceptual information, or from both.

Keywords: Spatial cognition; Embodied cognition; Geographical structures; Cognitive maps; Mental representations; Symbol interdependency; Latent semantic analysis

Cognitive maps are mental representations that contain spatial information about locations, as well as the directional relationships and distances between them (Downs & Stea, 1977; Tolman, 1948). These mental representations are formed through perceptual experiences, for example, via locomotion and stationary viewing, through static pictorial representations, such as diagrams, paintings, and photos, and via dynamic pictorial

Correspondence should be sent to Max M. Louwerse, Department of Psychology/Institute for Intelligent Systems, University of Memphis, 202 Psychology Building, Memphis, TN 38152. E-mail: maxlouwerse@gmail.com.

representations, including animations, movies, and videos (Bryant, Lanca, & Tversky, 1995; Bryant & Tversky, 1999; Friedman & Montello, 2006; Richardson, Montello, & Hegarty, 1999; Thorndyke & Hayes-Roth, 1982). Cognitive maps can also be formed through language, by means of descriptions of locations from the navigator's perspective, descriptions of tours of the environment, and descriptions of locations using a fixed point of view such as wind directions (Taylor & Tversky, 1996).

Various studies have demonstrated a functional equivalence for cognitive maps derived from spatial language and those derived from non-linguistic sources (Avraamides, 2003; Avraamides, Loomis, Klatzky, & Golledge, 2004; De Vega & Rodrigo, 2001; Franklin & Tversky, 1990; Loomis, Lippa, Klatzky, & Golledge, 2002). However, most studies on the relation between space and language have concentrated on local-scale locations (Ehrlich & Koster, 1983; Ehrlich & Johnson-Laird, 1982; Freundsuh & Mercer, 1995; Morrow, Greenspan, & Bower, 1987; Taylor & Tversky, 1992a,b; Tversky, Lee, & Mainwaring, 1999). Moreover, these studies typically focus on explicit spatial cues, such as the use of spatial prepositions in language or other explicit spatial cues (Coventry & Garrod, 2004; Taylor & Tversky, 1996). Relatively little is known about whether language encodes spatial information implicitly, for instance through its linguistic structure. Furthermore, it is not clear whether humans can form global-scale cognitive maps, such as locations of cities in countries or continents, using such implicit linguistic information. Finally, it is an open question to what extent human performance on estimating spatial information compares between a perceptual source of information (maps) and an implicit linguistic source of information (language) (Montello & Freundsuh, 1995).

How could the linguistic structure reveal spatial information, such as the location of cities? The answer might lie in semantic context. Firth (1957, p. 11) stated, "You shall know a word by the company it keeps" (Firth, 1957, p. 11). Analogously, you shall know the physical distance between locations by the lexical company they keep. Louwerse, Cai, Hu, Ventura, and Jeuniaux (2006); Louwerse, Hutchinson, and Cai (2012); and Louwerse and Zwaan (2009) tested this adage by predicting that cities that are located together are debated together. Using statistical linguistic frequencies, Louwerse et al. (2006) demonstrated that the semantic relations between cities in France have a negative correlation with the physical distances between those cities. Louwerse et al. (2012) tested this by predicting the locations of cities in China using Chinese text, and the locations of cities in the Middle East using Arabic text. Similarly, Louwerse and Zwaan (2009) took the largest cities in the United States and computed the semantic relations between them using three newspapers (*New York Times*, *Wall Street Journal*, and *Los Angeles Times*). The computation of semantic relations correlated with the longitude and latitude of the 50 cities, suggesting that language encodes spatial information.

The question whether cognitive maps are formed through language statistics or perceptual input is analogous to the debate whether cognition is symbolic or embodied (De Vega, Glenberg, & Graesser, 2008). The central question in this debate is whether conceptual processing predominantly depends on statistical linguistic structure or on perceptual simulations. Various studies have now explicitly acknowledged that the answer is

that conceptual processing relies on both (Barsalou, Santos, Simmons, & Wilson, 2008, Louwerse, 2007, 2011a,b, Louwerse & Jeuniaux, 2008). That is, statistical linguistic patterns and perceptual simulations interact with one another. For instance, Louwerse (2011b) proposed the Symbol Interdependency Hypothesis, arguing that language encodes embodied and perceptual relations, encodings that language users can use as a shortcut in conceptual processing. Various experiments comparing the effect of statistical linguistic frequencies and perceptual simulation on processing time have demonstrated that participants rely on both linguistic and non-linguistic factors (Louwerse, 2008), and in both language processing and picture processing (Louwerse & Jeuniaux, 2010), whereby less precise statistical linguistic processes precede precise perceptual simulation processes (Louwerse & Connell, 2011). However, these studies focused on conceptual processing, such as understanding the relation between the concepts *sky* and *ground*. It remains an unanswered question as to whether these findings can be extended to spatial cognition, specifically global-scale cognitive maps.

As we have discussed earlier, cognitive maps could have been formed over time from a person's exposure to language, from looking at maps, through locomotion, and many other sources (Louwerse & Zwaan, 2009; Montello & Freundsuh, 1995). It is therefore very difficult to disentangle effects that could be attributed to language statistics and effects that could be attributed to perceptual information for those spatial structures that humans are familiar with. A solution to this problem lies in selecting locations participants are likely unfamiliar with, for instance by using cities in an unfamiliar country (cf. Radvansky, Copeland, & Zwaan, 2005). That is, a novel could be selected whose story includes fictional cities located in a fictional world. An experiment could then compare performance of participants deriving spatial structure from the text versus from a map. With this solution, two practical problems emerge. First, what text is large enough to contain enough city locations allowing for an investigation of global-scale cognitive maps? Secondly, if such a text is found, how can participants be asked to read the same large amount of text that is desirable for computational linguistic analyses, while having a text that is small enough to control time on task (an in-lab experiment could not last more than a few hours)?

The fictional world as depicted by J.R.R. Tolkien in his map of Middle Earth and described as the setting in *The Hobbit* and the *Lord of the Rings* trilogy (LOTR) offers a solution. First, the LOTR texts consist of approximately half a million words, its size making it adequate for computational linguistic analyses. Secondly, the texts consist of 32 cities, making it suitable for a sufficient comparison of locations. Thirdly, the texts allow for a direct relation between the fictional text and the fictional map, both being based on the author's cognitive map. Yet the texts allow for a study comparing human text-based and map-based performance, because participant populations have already read the relatively large text.

In two studies, we investigated whether the language of LOTR encodes the relative spatial distances between the cities of Middle Earth (Study 1), and to what extent participant performance on estimating the city locations differed when participants primarily relied on the Tolkien text or a map of Tolkien's world (Study 2).

1. Study 1: Computational text-based spatial estimates of cities in Middle Earth

The first study tested the hypothesis that statistical linguistic frequencies encode spatial structure. One could investigate this hypothesis by computing first-order co-occurrences of city names (cf. Louwerse, 2011b; Louwerse & Zwaan, 2009). However, sparsity problems frequently emerge with first-order co-occurrence computations (Landauer, McNamara, Dennis, & Kintsch, 2007). That is, an extremely large corpus size is required to allow for two lexical items to co-occur in the first place (Louwerse, 2011b). Latent semantic analysis (LSA) provides a solution to this sparsity problem by not only relying on first-order co-occurrences but also higher-order co-occurrences mapping initially meaningless words into a continuous high dimensional semantic space (Landauer et al., 2007).

In this study, the input was the electronic version of the LOTR texts, the result of transcribing the pages of the texts. This electronic version was solely used for the research purposes mentioned here. The electronic document consisted of a total number of 629,570 words, segmented into paragraphs, from which a large term-document was created. For instance, if there are m terms in n paragraphs, a matrix of $A = (f_{ij} \times G(j) \times L(i, j))_{m \times n}$ is obtained, in this case $m = 16,140$ and $n = 11,659$. The value of f_{ij} is a function of the integer that represents the number of times term i appears in document j ; $L(i; j)$ is a local weighting of term i in document j ; and $G(j)$ is the global weighting for term j . The matrix of A has, however, lots of redundant information. Singular value decomposition reduces this noise by decomposing the matrix A into three matrices $A = U\Sigma V'$, where U is an m by m matrix and V is an n by n square matrix, with Σ being an m by n diagonal matrix with singular values on the diagonal. By removing dimensions corresponding to smaller singular values, the representation of each word is reduced as a smaller vector with each word now becoming a weighted vector on 300 dimensions, with only the most important dimensions that correspond to larger singular values being kept (Landauer et al., 2007). The semantic relationship between the words (cities) was estimated by taking the cosine between two vectors. We hypothesized that this cosine has an inverse relationship with the physical distance between the cities as presented on a map of Middle Earth.

1.1. Results and discussion

All 32 city names in Middle Earth were used for a 32×32 cosine matrix. Only cities were selected to avoid confusion regarding the exact location and names of rivers, mountain ridges, forests, and ruins (Table 1), even though such non-point locations will of course also contribute to spatial structure. The 32×32 cosine matrix was compared with the 32×32 distance matrix, obtained from the 2D coordinates of the cities on a map of Middle Earth (Fonstad, 2001; Fig. 1A).

There are two commonly used methods that compare the mapping of two planes. One method uses Multidimensional Scaling (MDS) to obtain two dimensions and applies bidimensional regression analyses to compare the coordinates. Whereas in a unidimensional regression each data point is shifted by intercept and slope, in bidimensional regression each actual and predicted value of the dependent variable is presented by a point in

Table 1
 Cities from *The Lord of the Rings* and *The Hobbit* being used in the analysis

Annúminas	Hobbiton
Archet	Isengard
Barad-dûr	Linhir
Bree	Longbottom
Bucklebury	Michel Delving
Bywater	Minas Anor
Calembel	Osgiliath
Caras Galadhon	Pelargir
Crickhollow	Rivendell
Combe	Scary
Dale	Staddle
Dol Amroth	Stock
Edoras	Tharbad
Frogmorton	Tuckborough
Mitholond	Waymeet
Haysend	Woodhall

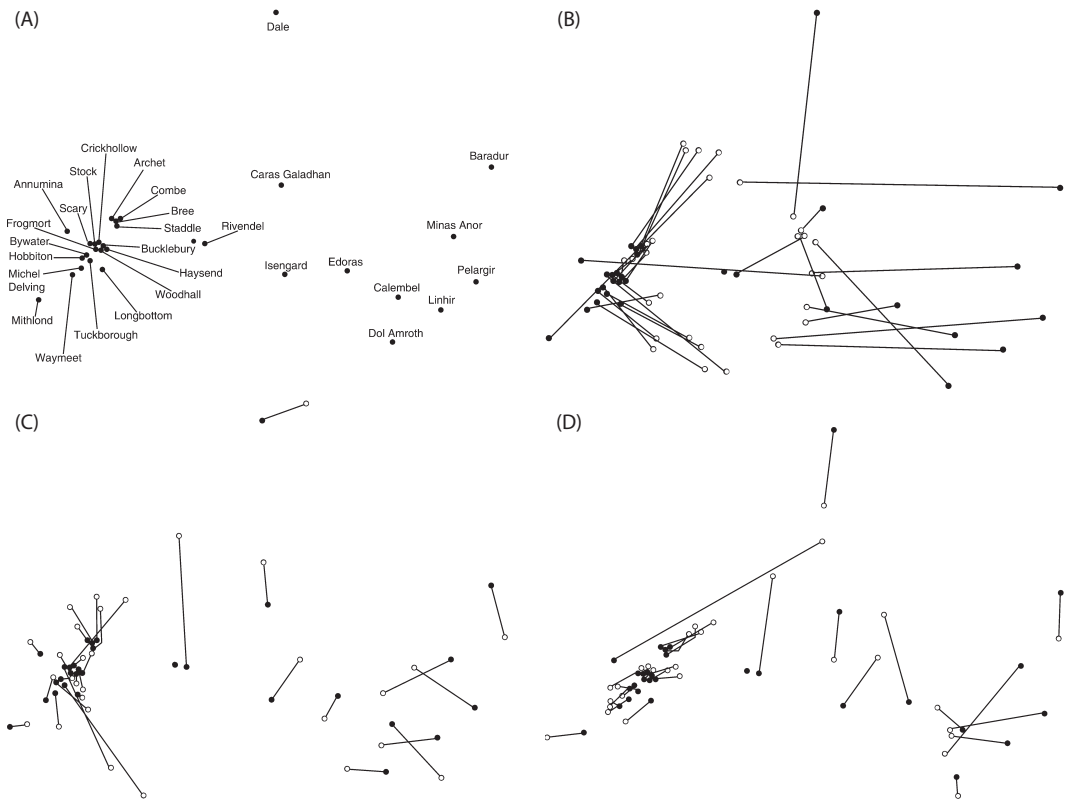


Fig. 1. Authentic Middle Earth map with 32 cities (A), computational estimates extracted from text (Study 1) (B), participants' estimates after studying map (Study 2) (C), and participants' estimates after studying text (Study 2) (D). Black circles mark authentic locations, white circles mark estimates, with the differences marked by a line.

space, whereby vectors represent intercept and slope (Friedman & Kohler, 2003; Tobler, 1994).

However, MDS relies exclusively on estimations of distances to obtain configural representations of locations and ignores directional information (Waller & Haun, 2003). A technique that considers both distance and directional information is the Procrustes technique (Schönemann & Carroll, 1970). This technique overlays the maps with the actual one and minimizes the differences between the maps by Euclidean transformations, such as translation, rotation, and scaling (Dryden & Mardia, 1988). We applied both bidimensional regression analyses and Procrustean analyses. The two methods should, however, yield equivalent results, since one maximizes least square fits (MDS, bidimensional regression) and the other minimizes the sum of squared errors (Procrustean analysis).

For the first analysis, the cosine matrix of 32×32 cosine values between LOTR city names was submitted to an MDS analysis. The MDS analysis was run using the ALSCAL algorithm (Young & Harris, 1990). A Euclidean distance measure transformed the semantic similarities into dissimilarities, such that the higher the value, the longer the distance. Default MDS criteria were used with an S-stress convergence of 0.001, a minimum stress value of 0.005, and a maximum of 30 iterations. We chose a low dimensionality to rule out overfitting of the data (Borg & Groenen, 2005). A two-dimensional scaling was moderate, with a Stress value = 0.36 and $R^2 = .46$. A bidimensional regression of the loadings of the first two MDS dimensions and the authentic longitude and latitude of the cities yielded a significant moderate correlation, $r = .46$, $p < .001$, $n = 32^1$ (Fig. 1B).

By comparison, a Monte Carlo simulation with 1,000 random reshufflings (and bidimensional regression analyses) of the 32 city coordinates revealed that the probability of random coordinate pairs having a correlation of $r = .46$ was <1 in 1,000 ($Mean\ r = .16$, $SD = 0.08$).

In an analysis alternative to the bidimensional regression analysis, the LSA cosine structure of 32×32 cities was computed and orthogonally Procrustes rotated toward the authentic structure of 32×32 city distances. This estimate also yielded a moderate significant correlation with the authentic locations of the cities in Middle Earth, $r = .57$, $p < .001$.

As before, we ran a Monte Carlo simulation of the Procrustes analysis with 1,000 random matrix reshufflings. The results again showed the probability of obtaining such a correlation was <1 in 1,000 ($Mean\ r = .22$, $SD = 0.07$).

These findings provide evidence that the LOTR texts encode spatial structure of cities in Middle Earth. Even though the text is relatively small for an LSA analysis (Landauer et al., 2007) and the correlations were moderate explaining 21–33% of the variance, LSA predictions using MDS and bidimensional regressions, as well as Procrustes analyses, demonstrated a significant relation between computational text-based estimates and authentic locations of cities in Middle Earth.

These findings replicate the findings by Louwerse et al. (2006) and Louwerse and Zwaan (2009), and show that the texts of *Lord of the Rings* and *The Hobbit* encodes the spatial structure of Middle Earth.

2. Study 2: Human spatial estimates of cities in Middle Earth

Study 2 investigated to what extent participants were able to estimate LOTR city locations after having read the text or after having studied a physical map of Middle Earth.

2.1. Method

2.1.1. Participants

Thirty-seven participants from the University of Memphis participated in this study for course credit or monetary compensation. Participants took a 20-item questionnaire consisting of detailed open-ended questions that were directly related to the contents of the LOTR texts, with five questions for each of the four books (see http://madresearchlab.org/materials/LOTR/LOTR_Questions.pdf). The purpose of this questionnaire was to determine whether participants had read the LOTR texts. Twenty-five participants with <20% accuracy were selected for the map-based condition. Twelve participants with more than 80% accuracy were selected for the text-based condition. Participants were not included in either condition on the basis of having seen these movies, because little spatial information could be extracted from the movies.²

2.1.2. Materials

The same 32 cities in Middle Earth featured in Tolkien's the LOTR texts were selected as the ones used in Study 1.

2.1.3. Procedure

The 25 map-based participants were given a sheet of paper with the cities in alphabetical order and a detailed physical map with red dots marking cities and numbers next to the dots corresponding with city names on a list. This map contained the contours of Middle Earth, rivers, mountain ridges, and forests. Participants were instructed to study the map for 20 min. The text-based participants were not given a map, and not given any study time.

Both text-based and map-based groups were then given a blank piece of paper with an x - and y -axis through the center, and the list of cities. Participants, were told to mark all 32 cities with the corresponding city number and/or dot. There were no restrictions on what participants could draw on the blank sheet and no time limit was enforced on the task. The entire study lasted no more than 1 h.

Both text-based and map-based participants were given access to a text-only electronic version of LOTR. This was important for the text-based participants, because they had not been informed about the purposes of the study. Map-based participants were also allowed access to the text to ensure that both groups received a similar amount of information.

2.2. Results and discussion

Participant-marked sheets were scanned into an electronic format. Pixel coordinates were identified for each of the 32 cities as marked by the participants. We conducted two series

of analyses. In the first series, we compared the participant estimates with the authentic locations of the cities in Middle Earth. In the second series, we compared the estimates with the computational estimates from Study 1. As before, we used both bidimensional regression and Procrustes analyses.

Bidimensional regression coefficients showed a significant correlation between the map-based estimates and the authentic location of the 32 cities, Mean $r = .82$ ($SD = 0.20$), $p < .001$, $n = 32$ (Fig. 1C). For the text-based estimates the correlation between participant estimates and the authentic locations was equally high, $r = .78$ ($SD = 0.11$) (Fig. 1D).

As before, we assessed the significance of symmetry using a symmetric Procrustes rotation whereby the participant structure was computed and orthogonally Procrustes-rotated toward the authentic structure. Map-based estimates yielded a correlation with the authentic coordinates, $r = .81$ ($SD = 0.18$), $p < .001$, $n = 32$. A similar result was obtained for the correlation between text-based estimates and authentic coordinates, $r = .77$ ($SD = 0.13$), $p < .001$, $n = 32$.

The difference between the performance of map-based and text-based coordinates was negligible. Perhaps this is no surprise: Both a linguistic source (text) and a non-linguistic source (map) can help to build a situation model or cognitive map (Montello & Freundschuh, 1995; Louwerse & Zwaan, 2009; Van Dijk & Kintsch, 1983). The important question here is whether the computational text-based estimates better predicted the participant text-based estimates than the participant map-based estimates. To answer this question, we compared the computational estimates with the text-based and the map-based human estimates, both in a bidimensional regression analysis and in a Procrustes analysis.

The bidimensional regression analysis for the map-based estimates and the LSA estimates yielded a significant correlation, $r = .42$ ($SD = 0.10$). For the text-based estimates this correlation was slightly higher, $r = .48$ ($SD = 0.07$).

A Procrustes rotation of the configurations also yielded a significant result for the map-based estimates and the LSA estimates, $r = .36$ ($SD = 0.13$), $p < .001$, as well as the text-based estimates and the LSA estimates, $r = .39$ ($SD = 0.04$), $p < .001$, again slightly higher for the text-based estimates.

To compare the map-based and text-based findings, and the authentic and LSA estimates, a mixed-effects model was run on the regression coefficients of the two studies. The model was fitted using the restricted maximum likelihood estimation (REML). F -test denominator degrees of freedom were estimated using the Kenward–Roger's degrees of freedom adjustment.

A significant difference was obtained between the map- and text-based estimates and the authentic coordinates, with the correlations between the participants' estimates and the authentic coordinates being significantly higher ($M = 0.81$, $SD = 0.03$) than the participants' estimates and the LSA estimates ($M = 0.44$, $SD = 0.03$), $F(1, 71) = 129.69$, $p < .001$. No difference was obtained between text-based and map-based estimates, $F(1, 71) = 0.07$, $p = .79$, showing that map-based and text-based estimates were on par. Interestingly, a significant interaction was obtained, $F(3, 70) = 44.80$, $p < .001$ (Fig. 2), showing that map-based estimates were more accurate than text-based estimates when these estimates were compared with authentic map locations, but text-based estimates were more

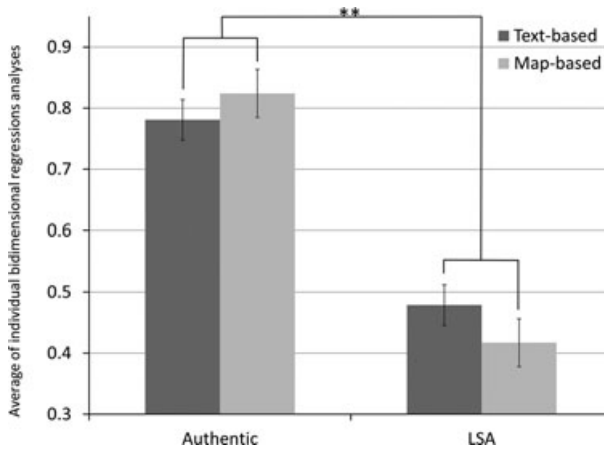


Fig. 2. Average r coefficients of individual bidimensional regression analyses. Asterisks mark significance ($p < .01$).

accurate than map-based estimates when these estimates were compared with estimates obtained from text.

3. General discussion

This study investigated whether the spatial structure of Middle Earth can be predicted using the *Lord of the Rings* and *The Hobbit* texts, and how human spatial estimates compare when these estimates primarily come from non-linguistic information (map of Middle Earth) versus linguistic information (text of LOTR). The results suggest that the physical distance between locations can be estimated by the lexical company they keep, leading to the conclusion that language encodes spatial structure. In a study where we compared human spatial estimates when those estimates were obtained from a map of Middle Earth or from the language of LOTR, we found no performance differences. The representations of spatial layouts derived from perceptual input are all functionally equivalent to the representations from linguistic input (Avraamides et al., 2004). This might suggest that linguistic input is immediately transformed in perceptual simulations, allowing the source for the cognitive map to be different (linguistic text or perceptual map) but the product of the cognitive map to be the same (perceptual map).

The difference between the accuracy of the text-based and map-based human estimates was small. This is not a surprise because cognitive maps can be formed both by using language and maps. Interestingly, however, the estimates obtained through the text-based statistical linguistic frequencies better predicted the human text-based estimates than the human map-based estimates. This result shows that even though human text-based and map-based performance might be on par, a difference can be observed when the source of information is taken into account.

Knowledge of words and knowledge of the world must somehow be linked (Malt & Wolff, 2010). The conclusion that performance on spatial tasks such as the one described in this study is mediated by linguistic and non-linguistic sources of information is very much in line of several theories of conceptual processing. For instance, Paivio's (1978, 1986) dual coding theory emphasizes that verbal and imagery processes interact. That is, verbal language of the LOTR text referentially evoked nonverbal mental images of the locations, events, and characters in the story (Sadoski & Paivio, 2001). Similarly, Barsalou et al.'s (2008) Language and Situated Simulation Theory states that the linguistic system and the simulation system are intrinsically related. These theories differ in one important aspect from the Symbol Interdependency Hypothesis discussed earlier, in that the latter assumes language encodes perceptual relations; these cues the language user applies in the comprehension process (Louwerse, 2007, 2011b). In previous work, we have also demonstrated that the linguistic system offers a quick and shallow heuristic that can provide good enough performance in certain tasks without recourse to deeper conceptual processing in a perceptual simulation system (Louwerse & Connell, 2011). Moreover, whether participants rely relatively more on linguistic or perceptual information depends on the cognitive task and the nature of the stimuli (Louwerse & Jeuniaux, 2010).

Even though it intuitively seems more likely that people rely on perceptual information in spatial thought, the results in this study suggest that linguistic statistical distributions can be considered to be a good, although less explicit, source for spatial information allowing to reconstruct Middle Earth through the statistical linguistic frequencies of the locations in *Lord of the Rings*.

Notes

1. In this analysis we avoided overfitting of the data. However, a moderate two-dimensional scaling (Stress value = 0.355 and an $R^2 = .457$) might suggest an underfitting of the data. To determine whether a better fitting would yield a higher correlation, we ran MDS analyses for 3–6 dimensions. Even though Stress values and R^2 values were obviously higher (3D: Stress = .25, $R^2 = .60$; 4D: Stress = .20, $R^2 = .69$; 5D: Stress = .16, $R^2 = .74$; 6D: Stress = .16, $R^2 = .74$), bidimensional regression values for the first two dimensions and the authentic longitude and latitude were not (3D: $r = .22$, 4D: $r = .41$, 5D: $r = .23$; 6D: $r = .24$). That is, when the first two dimensions of the 3–6 dimensional fittings were compared with the authentic locations, correlations ranged between $r = .23$ and .41. When all combinations of dimensions were compared (e.g., dimension 2 with dimension 5), correlations ranged between $r = .17$ and .49. This suggests that even with a moderate fitting of the MDS solution, the results are strongest for the first two dimensions (see also Louwerse & Zwaan, 2009).
2. A map of Middle Earth is printed in the LOTR books and is briefly presented in the LOTR movies. Perhaps this reduces the validity of keeping linguistic and non-linguistic sources apart. However, across LOTR editions, the specificity of the Middle Earth map differs. Across 11 LOTR editions we examined, less than a third of the cities

were displayed on maps, $M = 10$ ($SD = 7.51$). Furthermore, in the LOTR movies only eight of the 32 cities were mentioned in dialog, and a map of Middle Earth with city names was presented a total of three times for <30 s each time across all movies. When a map was shown, only eight cities could be distinguished. We therefore consider it unlikely that city locations are primarily obtained from the map in the books or the movies. At the very least, without the availability of a better alternative, the benefits of (a) the size of the corpus, (b) the access to a fictional world, and (c) the number of cities being included outweigh these drawbacks.

References

- Avraamides, M. N. (2003). Spatial updating of environments described in texts. *Cognitive Psychology*, *47*, 402–431.
- Avraamides, M. N., Loomis, J. M., Klatzky, R. L., & Golledge, R. G. (2004). Functional equivalence of spatial representations derived from vision and language: Evidence from allocentric judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 801–814.
- Barsalou, L. W., Santos, A., Simmons, W. K., & Wilson, C. D. (2008). Language and simulation in conceptual processing. In M. de Vega, A. M. Glenberg & A. C. Graesser (Eds.), *Symbols, embodiment, and meaning* (pp. 245–283). Oxford, England: Oxford University Press.
- Borg, I., & Groenen, P. J. F. (2005). *Modern multidimensional scaling: Theory and applications*. New York: Springer.
- Bryant, D. J., Lanca, M., & Tversky, B. (1995). Spatial concepts and perception of physical and diagrammed scenes. *Perceptual and Motor Skills*, *81*, 531–546.
- Bryant, D. J., & Tversky, B. (1999). Mental representations of spatial relations from diagrams and models. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *25*, 137–156.
- Coventry, K. R., & Garrod, S. C. (2004). *Saying, seeing and acting: The psychological semantics of spatial prepositions*. New York: Psychology Press.
- De Vega, M., Glenberg, A. M., & Graesser, A. C. (Eds.) (2008). *Symbols and embodiment: Debates on meaning and cognition*. Oxford, England: Oxford University Press.
- De Vega, M., & Rodrigo, M. J. (2001). Updating mental frameworks mediated by pointing and labeling under physical and imaginary rotation. *European Journal of Cognitive Psychology*, *13*, 369–393.
- Downs, R. M., & Stea, D. (1977). *Maps in minds: Reflections on cognitive mapping*. New York: Harper and Row.
- Dryden, I. L., & Mardia, K. V. (1988). *Statistical shape analysis*. Chichester, England: John Wiley & Sons.
- Ehrich, V., & Koster, C. (1983). Discourse organization and sentence form: The structure of room descriptions in Dutch. *Discourse Processes*, *6*, 169–195.
- Ehrlich, K., & Johnson-Laird, P. N. (1982). Spatial descriptions and referential continuity. *Journal of Verbal Learning and Verbal Behavior*, *21*, 296–306.
- Firth, J. R. (1957). *Papers in linguistics 1934–1951*. London: Oxford University Press.
- Fonstad, K. W. (2001). *The atlas of Middle-Earth*. Boston: Houghton Mifflin.
- Franklin, N., & Tversky, B. (1990). Searching imagined environments. *Journal of Experimental Psychology: General*, *119*, 63–76.
- Freundschuh, S. M., & Mercer, D. (1995). Spatial cognitive representations of story worlds acquired from maps and narrative. *Geographical Systems*, *2*, 217–233.
- Friedman, A., & Kohler, B. (2003). Bidimensional regression: A method for assessing the configural similarity of cognitive maps and other two-dimensional data. *Psychological Methods*, *8*, 468–491.

- Friedman, A., & Montello, D. R. (2006). Global-scale location and distance estimates: Common representations and strategies in absolute and relative judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 333–346.
- Landauer, T., McNamara, D. S., Dennis, S., & Kintsch, W. (Eds.) (2007). *Handbook of latent semantic analysis*. Mahwah, NJ: Erlbaum.
- Loomis, J. M., Lippa, Y., Klatzky, R. L., & Golledge, R. G. (2002). Spatial updating of locations specified by 3-D sound and spatial language. *Journal of Experimental Psychology: Human Learning, Memory, and Cognition*, *28*, 335–345.
- Louwerse, M. M. (2008). Embodied representations are encoded in language. *Psychonomic Bulletin and Review*, *15*, 838–844.
- Louwerse, M. M. (2011a). Stormy seas and cloudy skies: Conceptual processing is (still) linguistic and perceptual. *Frontiers in Psychology: Cognition*, *2*, 1–4.
- Louwerse, M. M. (2011b). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, *3*, 273–302.
- Louwerse, M. M., Cai, Z., Hu, X., Ventura, M., & Jeuniaux, P. (2006). Cognitively inspired natural-language based knowledge representations: Further explorations of latent semantic analysis. *International Journal of Artificial Intelligence Tools*, *15*, 1021–1039.
- Louwerse, M. M., & Connell, L. (2011). A taste of words: Linguistic context and perceptual simulation predict the modality of words. *Cognitive Science*, *35*, 381–398.
- Louwerse, M. M., Hutchinson, S., & Cai, Z. (2012). The Chinese route argument: Predicting the longitude and latitude of cities in China and the Middle East using statistical linguistic frequencies. *Proceedings of the 34th Annual Conference of the Cognitive Science Society* (pp. 695–700). Austin, TX: Cognitive Science Society.
- Louwerse, M. M., & Jeuniaux, P. (2008). Language comprehension is both embodied and symbolic. In M. de Vega, A. Glenberg, & A. C. Graesser (Eds.), *Symbols and embodiment: Debates on meaning and cognition* (pp. 309–326). Oxford, England: Oxford University Press.
- Louwerse, M. M., & Jeuniaux, P. (2010). The linguistic and embodied nature of conceptual processing. *Cognition*, *114*, 96–104.
- Louwerse, M. M., & Zwaan, R. A. (2009). Language encodes geographical information. *Cognitive Science*, *33*, 51–73.
- Malt, B., & Wolff, P. (2010). *Words and the mind: How words capture human experience*. Oxford, England: Oxford University Press.
- Montello, D. R., & Freundschuh, S. M. (1995). Sources of spatial knowledge and their implications for GIS: An introduction. *Geographical Systems*, *2*, 169–176.
- Morrow, D. G., Greenspan, S. L., & Bower, G. H. (1987). Accessibility and situation models in narrative comprehension. *Journal of Memory and Language*, *26*, 165–187.
- Paivio, A. (1978). Imagery, language, and semantic memory. *International Journal of Psycholinguistics*, *5*, 31–47.
- Paivio, A. (1986). *Mental representations: A dual coding approach*. Oxford, England: Oxford University Press.
- Radvansky, G. A., Copeland, D. E., & Zwaan, R. A. (2005). A novel study: Investigating the structure of narrative and autobiographical memories. *Memory*, *13*, 796–814.
- Richardson, A. E., Montello, D., & Hegarty, M. (1999). Spatial knowledge acquisition from maps, and from navigation in real and virtual environments. *Memory and Cognition*, *27*, 741–750.
- Sadoski, M., & Paivio, A. (2001). *Imagery and text: A dual coding theory of reading and writing*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Schönemann, P. H., & Carroll, R. M. (1970). Fitting one matrix to another under choice of a central dilation and a rigid motion. *Psychometrika*, *35*, 245–255.
- Taylor, H. A., & Tversky, B. (1992a). Descriptions and depictions of environments. *Memory and Cognition*, *20*, 483–496.
- Taylor, H. A., & Tversky, B. (1992b). Spatial mental models derived from survey and route descriptions. *Journal of Memory and Language*, *31*, 261–282.

- Taylor, H. A., & Tversky, B. (1996). Perspective in spatial descriptions. *Journal of Memory and Language*, 35, 371–391.
- Thorndyke, P. W., & Hayes-Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, 14, 560–589.
- Tobler, W. (1994). Bidimensional regression. *Geographical Analysis*, 26, 186–212.
- Tolkien, J. R. R. (2002/1937). *The hobbit*. Boston: Houghton Mifflin.
- Tolkien, J. R. R. (2002/1954). *The lord of the rings*. Boston: Houghton Mifflin.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55, 189–208.
- Tversky, B., Lee, P. U., & Mainwaring, S. (1999). Why speakers mix perspectives. *Journal of Spatial Cognition and Computation*, 1, 399–412.
- Van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension*. New York: Academic Press.
- Waller, D., & Haun, D. B. M. (2003). Scaling techniques for modeling directional knowledge. *Behavior Research Methods, Instruments, & Computers*, 35, 285–293.
- Young, F.W., & Harris, D.F. (1990). Multidimensional scaling: Procedure ALSCAL. In J.J. Norusis (Ed.), *SPSS base system user's guide* (pp. 396–461). Chicago, IL: SPSS Inc.

Appendix

There are two potential concerns that should be addressed. One concern in this study was that a map of Middle Earth was included in the LOTR books. Participants in the text condition who had read the LOTR could therefore have consulted the maps and have used the non-linguistic information instead of the text as a source for their estimates. To check for this possibility, the questionnaire also asked participants on a scale of 1–6 whether they had consulted the physical map while reading the book, 1 being never, 6 being frequently. The results showed participants sometimes consulted the map ($M = 2.5$, $SD = 1.6$). While the geographical knowledge of the participants in the text-based condition may have been influenced by non-linguistic information, the influence of non-linguistic information in the participants in text-based condition was probably far less than the influence of non-linguistic information in the participants in the map-based condition. However, to further investigate the effect of access to the map on the text-based estimates, we compared the eight participants who reported to have never or rarely consulted the map (four participants never consulted the map, four rarely consulted the map) and the participants who sometimes or frequently consulted the map (three participants sometimes consulted the map; one participant frequently consulted the map).

When the regression coefficients of the authentic coordinates and the text-based estimates were compared between participants who hardly consulted the map, no differences were found, $F(1, 11) = 2.17$, $p = .17$, $M = 0.75$ ($SD = 0.13$) and $M = 0.85$ ($SD = 0.04$), respectively. Similarly, no differences were found between the two participant groups for the regression coefficients obtained from computer estimates and human estimates, $F(1, 11) = 0.46$, $p = .51$, $M = 0.49$, $SD = 0.09$ and $M = 0.46$ ($SD = 0.02$), respectively.

Procrustes analyses showed similar findings. No difference was found between the two participant groups when the correlations with the authentic coordinates were compared, $F(1, 11) = 2.60$, $p = .14$, $M = 0.71$ ($SD = 0.20$) (less map consultation) and $M = 0.88$

($SD = 0.04$) (more map consultation) or the computer estimates were compared, $F(1, 11) = 0.45$, $p = .52$, $M = 0.51$ ($SD = 0.10$) (less map consultation) and $M = 0.48$ ($SD = 0.05$) (more map consultation).

If viewing the map more frequently did not impact the accuracy of the cognitive maps, map consultation cannot be considered a confounding variable (see also Note 2).

A second concern relates to consultation of the Tolkien text during the experiment. As a memory cue, we had provided participants with an electronic version of the text—after all, participants might not remember some of the city names out of context anymore. The spatial estimates might therefore have come from specific text searches during the experiment, instead of the co-occurrences of the cities. To rule out this possibility, we conducted a study in which we investigated to what extent access to the electronic version of LOTR affected the performance of participants. Twenty-five native speakers of English with no prior knowledge of LOTR participated in this study. They were given access to the electronic version of LOTR and asked to locate the 32 cities on a piece of paper. No difference was found between participants who had access to an electronic copy of the text and those who did not for regression coefficients based on a comparison with the authentic coordinates, $F(1, 49) = 1.66$, $p = .20$, $M = 0.73$ ($SD = 0.28$) (map only) and $M = 0.82$ ($SD = 0.19$) (map and access text). No difference between these two participant groups was found for the computational estimates either, $F(1, 49) = 0.42$, $p = .52$, $M = 0.40$ ($SD = 0.13$) (map only) and $M = 0.42$ ($SD = 0.10$) (map and access text). A similar result was obtained for the Procrustes analyses, $F(1, 49) = 1.56$, $p = .22$ (authentic coordinates) and $F(1, 49) = 0.86$, $p = .36$ (computational estimates).