

The Distance–Similarity Metaphor in Network-Display Spatializations

**Sara Irina Fabrikant, Daniel R. Montello,
Marco Ruocco, and Richard S. Middleton**

ABSTRACT: Dimensionality reduction algorithms are applied in the field of information visualization to generate low-dimensional, visuo-spatial displays of complex, multivariate databases—spatializations. Most popular dimensionality reduction algorithms project relatedness in data content among entities in an information space (e.g., semantic similarity) onto some form of distance among the entities, such that semantically similar documents are placed closer to one another than less similar ones. In previous studies of point-display spatializations we have shown that people indeed associate metric straight-line inter-point distances with the semantic dissimilarity of documents depicted as points in two-dimensional space. In this paper we investigate the strategies viewers employ when conflicting notions of distance (straight-line metric vs. network metric vs. topological proximity) are jointly shown in a spatialized network display of Reuters news articles depicted as points connected by links. We report empirical results of an experiment where viewers are asked to assess document similarity, depending on various distance types. We also investigate how cartographic symbolization principles (the use of visual variables, such as size, color hue, and value) influence similarity judgments. These findings provide rare empirical evidence for generally accepted design practices within the cartographic community (e.g., the effects of visual variables). In addition, empirical results from this and related studies can be used to develop design guidelines for constructing cognitively adequate spatializations for knowledge discovery in very large databases. We conclude by presenting design guidelines for network spatializations within the context of cartographic practice and theory.

KEYWORDS: Visualization, spatialization, perception, network, distance, similarity, visual variables

Introduction

Spatialization is the process of generating an information display of non-spatial data. Information spatialization is inspired by the intuition that spatial or graphical displays (e.g., maps, charts, photographs, diagrams) can help to amplify cognition (Tversky 2000), whether or not viewers are familiar with the construction details of the viewed display (Card et al. 1999). Generally, non-expert viewers do not know how spatializations are created and are not told through legends or other traditional map marginalia how to interpret the displays.

Spatializations typically rely on dimension reduction techniques (e.g., multidimensional scaling, self-organizing maps) and layout algorithms (e.g., energy minimization/force-feedback models) to project relatedness (e.g., similarity) in non-spatial data content onto distance, such that semantically similar documents are placed closer to one another than are less similar ones in an information space (Börner et al. 2003). We have coined this widely applied design principle the “distance–similarity metaphor” (Montello et al. 2003). In information

Sara Irina Fabrikant, Assistant Professor, **Daniel R. Montello**, Professor, **Marco Ruocco** and **Richard S. Middleton**, graduate student researchers, are all at the Department of Geography at the University of California, Santa Barbara, CA 93106, USA. Tel: (805) 893-5305, Fax: (805) 893-3146. Email: <{sara, montello, ruocco, middleton}@geog.ucsb.edu>.

visualizations (Chen 1999), especially those depicting knowledge-domains (Chen 2003), node-link displays are popular graphical devices for expressing the distance–similarity metaphor. A good example is the TouchGraph Google Browser¹, which provides a node-link display as a graphical user interface to Google’s “what’s related” search option, used to find web pages similar to a queried page.

Figure 1 depicts a network of web pages related to the highlighted web page labeled “Saraland” in the center of the display (the web page belongs to the first author). The placement of the nodes is achieved with a force-feedback graph-layout algorithm (e.g., Kamada and Kawai 1989). Nodes are re-arranged in a spring-like fashion, so that connected web sites not only get linked, but more strongly connected web sites contract towards each other in the spatialization. The resulting configuration is modified aesthetically,

¹ On the web at <http://www.touchgraph.com/TGGoogleBrowser.html>.

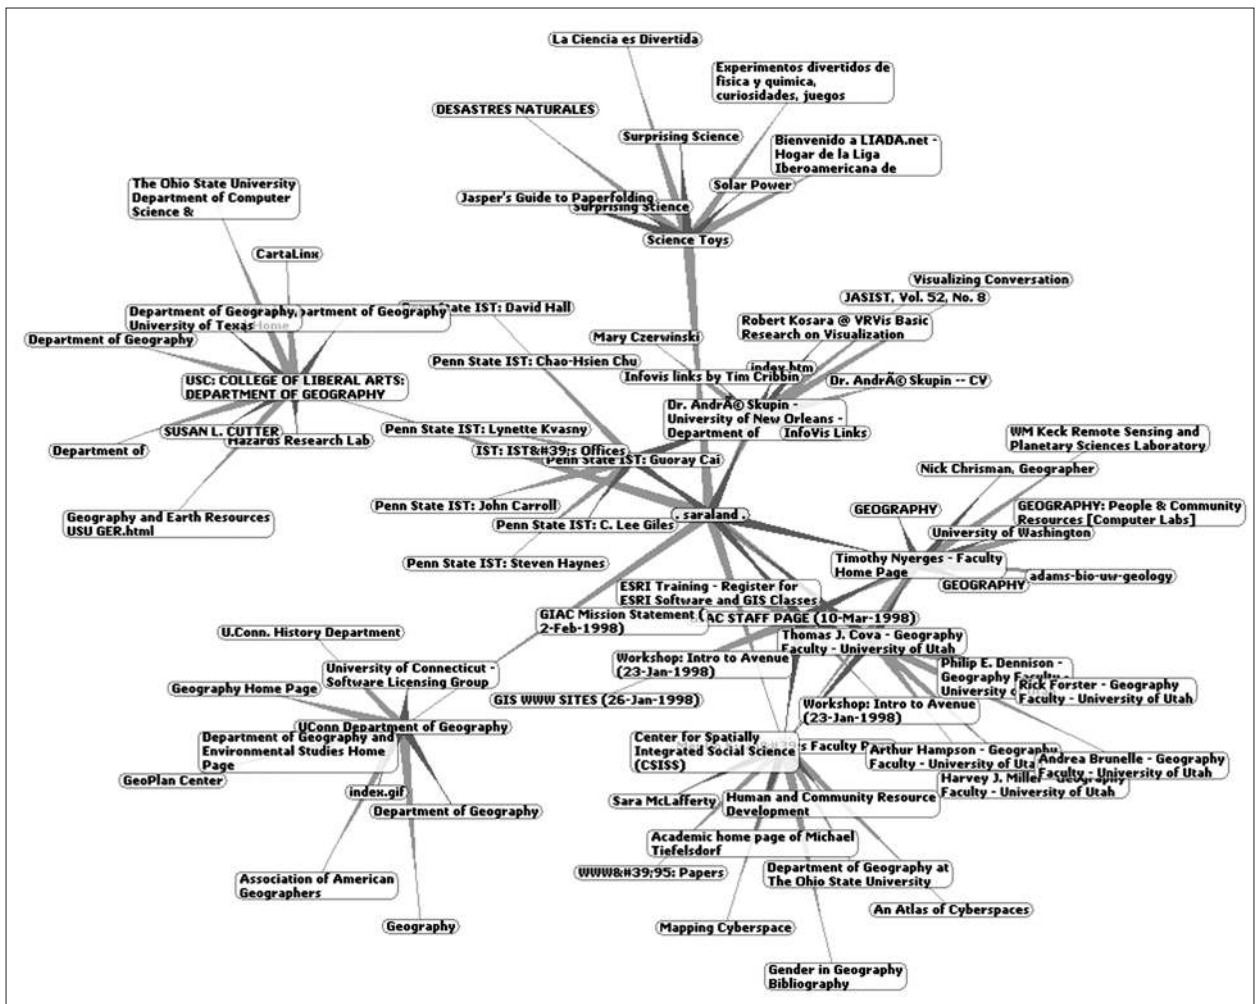


Figure 1. Network spatialization display depicting linkage relationships between the first author's web site and other sites on the web.

including reducing edge crossings to a minimum and evening edge lengths to achieve a balanced and visually pleasing layout.

Especially popular as a way to depict semantic relatedness has been the Pathfinder Network Scaling (PFNET) approach, derived from author co-citation analyses or word co-occurrences in text documents. Dominant relationships in proximity data are represented as a 2D or a 3D graph in PFNETs, and when depicted in 2D, the PFNET looks like a network map, similar to the one depicted in Figure 1. PFNET computes a pattern of links based on pair-wise assessment of proximity information associated with nodes in the graph. Although all the nodes are included in the PFNET, the network is typically computed with a minimum of links, such that only the strongest proximity relationships between nodes are depicted, as shown in Figure 2.

Figure 2 depicts a subset of a larger PFNET of Reuters news articles. The whole network contains 504 docu-

ments of randomly selected news stories, collected during February 9-10, 2000 (Fabrikant 2001). The nodes represent individual news articles and the links represent semantic relationships between the articles based on an analysis of word co-occurrence computed by the latent semantic analysis techniques (LSI) (Deerwester et al. 1990). The labels were extracted automatically from the article to indicate the main themes of the documents². In essence, a PFNET graph is akin to a minimum spanning tree. Which nodes get linked is determined by a chosen direct-distance metric, but unlike the minimum spanning tree, the triangle inequality requirement for metric spaces is relaxed (Schvaneveldt 1990).

Background

Point-display spatializations (e.g., multidimensional scaling plots) depict documents as unlinked points in space, semantic relatedness being reflected

² We would like to thank Dr. Bruce Rex at the Pacific Northwest National Laboratory for providing us with the Reuters news database.

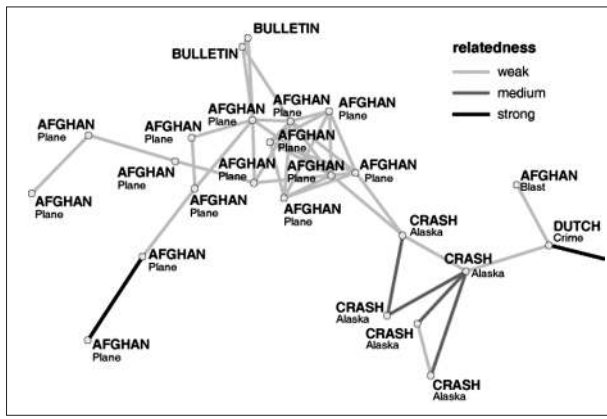


Figure 2. A force-directed graph depicting a Pathfinder Network of Reuters news stories.

in the proximity relationships of the points (Montello et al. 2003). The distance–similarity metaphor on a Pathfinder Network map, however, may operate at a variety of conceptual levels besides just straight-line or direct distance. Consider a geographic network transportation map such as the famous London Underground chart shown below in Figure 3 (network maps are discussed by Ruggles and Armstrong 1997). This network map depicts underground stops as nodes and the connections between them as links. Several distance types are jointly shown on one map display:

- Physical separation between locations measured along a straight-line, across the network (direct metric distance);
- Physical separation between locations measured along the network links (network metric distance); and
- Nonmetric measures of proximity, such as the number of nodes or stops between source and destination along the links (network topology).

When inspecting the map in Figure 3 it may not be obvious to a first-time visitor to London which Tube stop is closest to Cannon Street (labeled A in Figure 3)—Blackfriars (1) because it is closer when traveling on the network, or St. Paul’s (2) because it is closer in direct distance? If direct distance is not considered at all, do people use the number of stops between source and destination as a distance measure, or do they use the metric length on the map of the route between source and destination?

The London Underground map is a noteworthy example in the context of our study on network spatializations in many respects. First, it is considered a classic in information design history (Garland

1995). Mass transit maps depict a geographic network, but one in which the locations and distances between real-world places are typically systematically distorted, and are thus geographically inaccurate. Distance relationships on transportation maps are often distorted by the graphic designer in order to improve legibility of the graphic layout. The design emphasis is on the topology of the network, not the accuracy of the metric space. Often, angles are re-aligned to the cardinal directions (N-S-E-W), resulting in easily perceivable edge orientations of 0°, 45°, and 90° angles³. Transportation charts are a good example of attempts to achieve *cognitive adequacy* in graphic design. Their design reflects an implicit cognitive theory that travelers only need or want information about topology in order to get to destinations efficiently on subways. Also, cognitive research has confirmed that people tend to align the true orientation of geographic features along the orthogonal axes in their mental maps of real places (Montello 1991; Tversky 1981).

One of the tenets of information visualization is that spatialized displays work because they can be intuitively explored as if they represent real-world space (Wise 1999). Considering issues of cognitive adequacy in real-world network maps (e.g., the London Tube map), the research question arises as to whether the interpretation of metaphorical networks (such as Figures 1 and 2) may be different from the interpretation of real-world networks, more variable and less tied to assumptions about real space. Considering how plausible it is that people make judgment errors about the closeness of places depicted on subway maps—due to erroneous distance relationships introduced by the graphic designer to make a nice looking layout—how do people perceive semantic similarity expressed via graphic distance on network-display spatializations? As mentioned above, designers also systematically distort network spatializations, such as those shown in Figures 1 and 2, in order to generate what they believe to be cognitively adequate and aesthetically pleasing displays.

In previous studies, we have shown how the distance–similarity metaphor works in point- and surface-display spatializations (Fabrikant and Buttenfield 2001; Montello et al. 2003). People indeed associate metric graphical interpoint distances with the semantic similarity of text documents depicted in 2D and 3D. However, these studies also show that non-spatial visual variables, such as color hue and value, affect test outcomes for surface and point spatializations (Fabrikant 2003).

³An excellent map animation (Macromedia Flash) showing the distortion of the 2003 Tube map compared to the actual layout of the London underground is provided at <http://tube.tfl.gov.uk/content/tubemap/default.asp#flash>.

Presence of color improves the number of correct responses and can enhance or confuse the distance–similarity metaphor. Response times increased when the two Gestalt grouping principles of similarity and proximity were displayed jointly in a conflicting manner. Participants took additional time to decide how to resolve the visual conflict. If the color of the closest point/area is the same as of the reference point/area, response times are shorter and accuracy is increased. For surface spatializations, for example, applying color progressions based on the cartographic convention “darker is more” seems to increase the accuracy of responses and reduce response times. Interestingly, this provides an example where cartographic conventions even supersede environmental perceptual cues. When evaluating the metaphorical mapping of information density onto the spatial metaphor of clustering (as suggested by Wise 1999), surfaces in 2D and 3D that were rendered with a color range simulating natural terrain with “snow capped mountains” (white at the top of the color range) yielded less accurate responses than surfaces rendered with the inverted color range (Fabrikant 2003).

We are motivated by the general research question of how people decode the distance–similarity information embedded in network spatializations, and specifically what type of distance is equated with similarity, how it is equated, and how non-spatial visual variables (e.g., color, texture, and size) might influence people’s similarity judgments.

Experiments

In the present research, we conduct two experiments with non-expert users on the interpretation of the distance–similarity metaphor in node-link displays, similar to those shown in Figures 1 and 2. We call such information visualizations *network-display spatializations*. Each node represents an information-bearing entity such as a book, web site, or news story, and the links reveal the underlying semantic similarity between the information entities. In our experiments, we investigate how users interpret network-display spatializations to infer similarity relationships.

In different trials, participants make similarity judgments while viewing network displays that vary the spatial relationship between two pairs of comparison points (documents), the context provided by the network structure in the display, and the visual characteristics of the network links. We are specifically interested in the distance interpretation strategies viewers employ when conflicting notions of distance (e.g., direct distance vs. net-



Figure 3. A subset of the London Underground Map (© Transport for London).

work distance) are jointly shown in a spatialized network layout of information documents such as news articles. Our work also investigates how non-distance factors (i.e., non-spatial visual variables) may influence people’s similarity assessments. A third research question we pursue is how various types of distances may interact with non-distance factors such as the visual variables of size, color value, and hue that are commonly employed on maps. Our research is guided by geographic information theory, including cartographic design principles, and by psychological research on the perception and cognition of distance and similarity.

Experiment 1

In Experiment 1, we investigated how non-expert users interpret simple network spatializations, like those in PFNET maps. In addition to an initial exploration of how users interpret such displays, we wanted to test our data-collection methods, including our similarity comparison request and rating scale. We showed research participants computer displays of simple networks, explaining that the points at the nodes represented documents. Three of the document points were labeled ‘A,’ ‘1,’ and ‘2’; participants were asked to compare the similarity of A and 1 to the similarity of A and 2. A 9-point scale was provided for them to express their judgments of similarity on an interval scale. For the purpose of statistical analysis, we treated these scales as interval-level data so we could apply parametric techniques that are more powerful and flexible. This is commonly done in the psychometric literature, and it is justified in as much as such scales have been shown to approximate interval measurement rather than mere ordinal ranks (e.g., Howell 2002), even though they do not produce exactly equal intervals.




Figure 4. Sample screenshot from an Experiment 1 trial, showing display, similarity question, and rating scale as they appeared to participants.

Thus, the main task of the experiment required participants to compare the relative similarity of two pairs of document points, interconnected by network links. We varied the network displays across trials with respect to the distances between comparison points in three ways:

1. *Direct*—direct (or straight-line) metric distance between points, ignoring the network structure;
2. *Network*—network metric distance between points, measured along network links; and
3. *Node*—the number of nodes intervening between points, essentially a qualitative expression of proximity.

In addition to the network trials, this first experiment also tested other spatialization metaphors, namely points without networks and points within regions. Results from trials involving point spatializations are reported in Montello et al. (2003); results involving regions will be reported in a forthcoming paper. In the present paper, we focus exclusively on reporting results from the network displays.

Methods

Participants. Forty-four students (25 males and 19 females with a mean age of 21.0 years) from an undergraduate regional geography class took part in the experiment. They received a small amount of course credit in return for their participation. Given the background data collected in the pre-test questionnaire, the test population for both experiments was judged to be a good sample of the desired novice user population. The vast majority of the participants rated their map reading ability as average, have used maps only occasionally, and have never had training in cartography, GIS, computer graphics, or graphic design.

Materials. Participants viewed computer displays created using the Environmental Systems Research Institute’s (ESRI) ArcMap and composed of black points connected by blue network links (all points were on the network). The displays were inspired by PFNET displays such as that in Figure 2 but were not actual PFNET outputs. Each point was intended to represent a single document in a digital database. In each display, three points to be compared for similarity were labeled with red text as ‘A,’ ‘1,’ and ‘2’ (see Figure 4). Participants were prompted to “compare the similarity between document A and document 1 with the similarity between document A and document 2.” They rated similarity on a 9-point scale ranging left to right from ‘5’ to ‘1’ and then back up to ‘5.’ On the left, ‘5’ was labeled “Documents A and 1 are much more similar to each other.” In the middle, ‘1’ was labeled “1 and 2 are equally similar to A.” On the right, ‘5’ was labeled “Documents A and 2 are much more similar to each other.” In this paper, we refer to the pair of documents A and 1 as ‘A:1’ and A and 2 as ‘A:2.’

Participants viewed 10 different network trials in a block (as mentioned above, they also viewed 30 additional trials involving other display metaphors). The network displays were varied to allow comparisons of the effects of different distance relationships on judged similarity. Three different distance relationships were available in each display: Direct, Network, and Node. These were varied across trials in order to allow comparisons of the relative importance of each in judgments of similarity. Graphical elements that we did not expect to affect similarity judgments (such as the absolute location of the point on the screen) were varied non-systematically.

Participants were introduced to the concept of similarity, the style of the trials, and the format of the response scale through three practice trials at the beginning of the test. To avoid priming any particular equivalence between distance and similarity, the practice trials prompted judgments of non-distance similarity (by asking for a comparison of the similarity of images of, e.g., a pet dog, a domestic cat, and a tiger). Participants also responded to 11 pre-test questions about their personal backgrounds, including questions on age, gender, the presence of visual impairments (including specifically color blindness), as well as their formal experiences in such areas as cartography and GIS. After the main test questions, participants responded to 28 post-test questions that asked, for example, how useful they thought each display type was for rating similarity and how easy it was to judge similarities for each display type. Participants also indicated how they

had judged similarity and whether the displays reminded them of anything.

The experiment was administered using a Windows 2000 Pentium III personal computer. The interface was programmed with Microsoft Visual Basic 6.0. Images were projected onto a back-projection screen using an RGB color projector, generating an image size of 1.8 meters wide and 1.4 meters high, at 0.6 meters above the floor. Participants sat at a viewing table 2.7 meters in front of the screen, which enabled a horizontal viewing angle of slightly less than 40°. A standard mouse and keyboard were used to answer questions. Answers were recorded automatically and stored digitally, including the time required to make similarity judgments. Response time was measured as the elapsed time in milliseconds between the trial display appearing on the screen and the participant proceeding to the next trial.

Procedure. Participants were first told they would be presented with a series of trials about “diagrams that show an information collection from our computer database. The database contains documents such as news stories, books, and journal articles.” Participants were told that each document would be shown as a single point. No information was provided on how to judge similarity and no meaning was attached to the graphical elements other than the points. Participants were assured that there were no right or wrong answers, and they were asked not to waste time, as their answers would be timed.

Participants then answered the pre-test questions and performed the practice trials. The practice trials allowed participants to get comfortable with the test environment and gave them the opportunity to practice answering the questions by clicking on the desired number on the rating scale. Following that, participants responded to the main test trials organized into separate blocks (the block of network displays plus blocks of the other display types), so that participants rated all trials of one display type before turning to another type. Trials within each block were presented in a different randomized order for each participant. After completing the main test trials, participants answered the post-test questions, were marked down for credit, and thanked for their participation.

Results and Discussion

Similarity ratings were treated as 9-point interval scales, by scoring a response of ‘5’ to the far left (“A and 1 much more similar”) as a ‘1,’ a response of ‘5’ to the far right (“A and 2 much more similar”) as a ‘9,’ and a response of ‘1’ in the middle (“1 and 2 equally similar to A”) as a ‘5’ (see Figure 4). Thus




Figure 5. Sample display and similarity rating for trial of Experiment 1. All three types of distance are equal for A: 1 and A:2 in this trial; consequently, the two pairs of comparison documents were rated as equally similar.




Figure 6. Sample display and similarity rating for trial of Experiment 1. Direct and Network distances are equal for A:1 and A:2 in this trial, while Node distance is much greater for A:2 than A:1. The two pairs of comparison documents were again rated as equally similar.

a mean rating less than 5.0 indicates that participants saw A:1 as more similar, while a mean rating greater than 5.0 indicates they saw A:2 as more similar. Differences from equal similarity between A:1 and A:2 were thus tested with t-scores based on the difference of the mean similarity rating from 5.0.

An examination of similarity ratings for specific trials suggests the relevance of Network distance over Direct and Node distances. Figure 5 depicts a trial in which all three types of distance (Direct, Network, Node) were equal between A:1 and A:2. As expected, the mean similarity rating of 5.2 (sd = 1.6) was not significantly different from 5.0 for this trial ($t[43] = 0.75$).

Figure 6 depicts a trial in which Direct and Network distance were equal between A:1 and A:2, but Node




Figure 7. Sample display and similarity rating for trial of Experiment 1. Network and Node distances are greater for A:1 than for A:2, while Direct distance is greater for A:2. A: 2 is rated as significantly more similar than A:1.

distance was much greater for A:2 (four intervening nodes) than A:1 (one intervening node). Again, however, the mean similarity rating of 5.0 ($sd = 2.0$) was not significantly different from 5.0 for this trial ($t[43] = 0.16$). This suggests that participants did not interpret Node distance as relevant to similarity.

Finally, Figure 7 depicts a trial in which Direct distance was much greater (about twice) between A:2 than between A:1, but Network distance was greater for A:1 than A:2 (again about twice). The mean similarity rating of 6.5 ($sd = 2.1$) was significantly greater than 5.0 for this trial ($t[43] = 4.82$, $p < .001$). This shows that participants interpreted Network distance, and not Direct distance, as relevant to similarity. The pattern of results suggested by these three example trials—the influence of Network distance on similarity, and the lack of influence of Direct and Node distances—was supported by the results of all 10 trials of Experiment 1.

To interpret these data more systematically, we aggregated mean similarity judgments for each type of distance across trials that displayed the relationship of the distances A:1 and A:2 in the same way. That is, for each of the three types of distance, we compared the subset of trials (of 10) that equated the distances between A:1 and A:2 to the remaining subset of trials, in which the distance between A:1 was either less or greater than the distance between A:2. Similarity ratings for trials in which A:2 were closer were reverse-scored to allow aggregation with trials in which A:1 were closer. Table 1 presents mean similarity ratings for the trials aggregated into these two subsets, for each type of distance.

When A:1 and A:2 were the same distance apart for all three types of distance, mean similarity ratings for A:1 and A:2 were very close to a neutral rating of exactly 5.0. In the case of Direct distance, the mean

Type of Distance	Relationship of Distances A:1 and A:2	
	A:1=A:2	A:1<A:2
Network ^a	5.0	3.8***
Direct ^b	5.3**	5.1
Node ^c	5.2	6.1***

Notes. $N=44$ participants. ^aA:1=A:2 for 6 trials, A:1<A:2 for 4 trials. ^bA:1=A:2 for 7 trials, A:1<A:2 for 3 trials. ^cA:1=A:2 for 2 trials, A:1<A:2 for 8 trials. * $p < .05$, ** $p < .01$, *** $p < .001$, significance of difference from neutral similarity 5.0.

Table 1. Mean similarity ratings for trials in which network, direct, and node distances A:1 and A:2 equaled or differed, Experiment 1.

similarity of 5.3 was significantly different from 5.0, given the small standard deviation of this aggregated test, but this can be explained by the fact that in a couple of trials in which Direct distance was equal, Network distance was unequal. Conversely, in trials in which the Network distances between A:1 and A:2 were different, mean similarity was significantly greater for the closer pair of documents (by over a full scale point of similarity). No such difference for trials differing in Direct Distance was found. In trials in which A:1 and A:2 differed in Node distance, mean similarity was again significantly greater for the closer pair of documents (by just over a half-scale point of similarity). This apparently reflects the fact that Node distance was strongly correlated ($>.80$) with Network distance in our stimulus set. In trials where the Network and Node distance diverged (Figure 6 provides a good example), Network distance rather than Node distance clearly determined rated similarity.

We took another systematic approach to interpreting our data. The aggregated means presented above only address whether the three types of distance had any relation to similarity judgments. But one might expect the distance–similarity metaphor to imply that viewers of spatializations will quantitatively equate map distance with similarity (actually dissimilarity), at least approximately. If so, the degree of distance separation will be equivalent to the degree of similarity (i.e., twice a difference in distance will be interpreted as twice a difference in similarity). To evaluate this in the data, we calculated mean correlations of each of the three types of distance with rated similarity. Three Pearson's correlation coefficients were calculated separately for each participant, with the number of pairs of data points equal to the number of network trials (10 in this case). These correlations were strongest for Network distances: They were positive for 36 of the 44 participants, and averaged .49 across all participants (after transformation by Fisher's r -to- z). Based on a t -test calculated on the z -scores, this

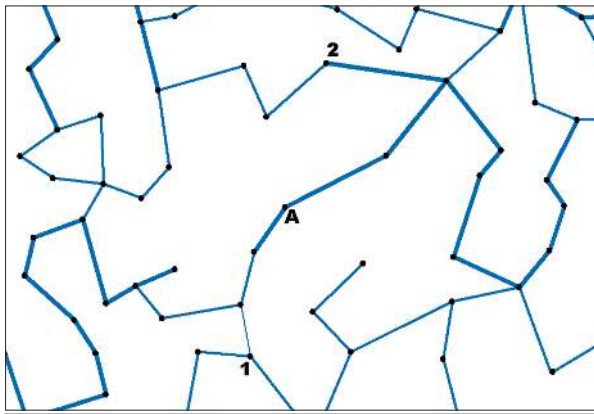


Figure 8. Sample display for a *width* trial of Experiment 2. A:1 is connected by mixed width links; A:2 is connected by all thick links.

correlation is significantly greater than 0, $t(43) = 6.76$, $p < .001$. The correlation of rated similarity with Direct distance was weak (.14). The correlation of rated similarity with Node distance was strong (.40), but this again reflects the fact that Node distance was inadvertently correlated with Network distance in our stimulus set.

Participants took a mean of 13.4 seconds (SD = 9.1 s) to respond to each network-display trial (the fastest mean response over all network trials was 5.3 s, the slowest was 56.8 s). Response-time varied significantly as a function of the order in which participants saw the block of network-display trials, $F(2, 41) = 4.79$, $p < .05$. Participants got faster as the experiment proceeded: Mean time-to-rate similarity for network displays in the first, second, and third block of trials was 17.9 s, 11.9 s, and 9.3 s, respectively.

Experiment 2

In our second experiment, we attempted to replicate and extend the results we obtained with network-display spatializations in our first experiment. First, we wanted to find out if our conclusion from Experiment 1 about the important role of Network distance (metric distance along network links) in similarity judgments, over Direct and Node distance, would be replicated with a new set of displays viewed by a new set of participants. Our comparison of Network-to-Node distances in Experiment 1, in particular, was confounded somewhat by the fact that these two types of distance were strongly correlated in our stimulus set.

We also wanted to examine the way three non-spatial visual variables might affect judgments of similarity in network displays, namely the width, value (lightness), and hue of links connecting comparison documents. These link variables may affect judgments of similarity in at least three,




Figure 9. Sample display for a *value* trial of Experiment 2. A:1 is connected by all dark links; A:2 is connected by all light links.

nonexclusive, ways. First, the absolute or relative level of the link variable may be semantically mapped onto similarity (e.g., dark or darker may be interpreted as more similar). A second possibility is that the level of the link variable may connote similarity, depending on the levels of other links or the surrounding context (e.g., darker may mean more similar only when it is embedded in a field of lighter links). A third possibility is that homogeneous link structures will connote greater similarity than heterogeneous link structures, no matter their absolute or relative levels (e.g., all light links connecting two documents indicates greater similarity than do links mixed in value).

Methods

Participants. Thirty-five students (18 males and 17 females with a mean age of 19.8 years) from an undergraduate introductory human geography class took part in the experiment. They received a small amount of course credit in return for their participation. None had participated in Experiment 1.

Materials. As in Experiment 1, participants viewed computer displays composed of different graphical elements. However, only network displays were tested in this experiment. All displays again included black points, with three points to be compared for similarity labeled as 'A,' '1,' and '2.' Participants again judged similarity, using the same scale as in Experiment 1. There were 65 network trials in this experiment, which were again varied systematically to allow comparisons of the effects of different visual variables on judgments of similarity.

Because only Network distance had been found relevant in Experiment 1, all trials in Experiment 2 varied Network distances between A:1 and A:2 but maintained equal Direct distances. Node distances




Figure 10. Sample display for a *hue* trial of Experiment 2. A:1 is connected by mixed hue links; A:2 is connected by all green links.

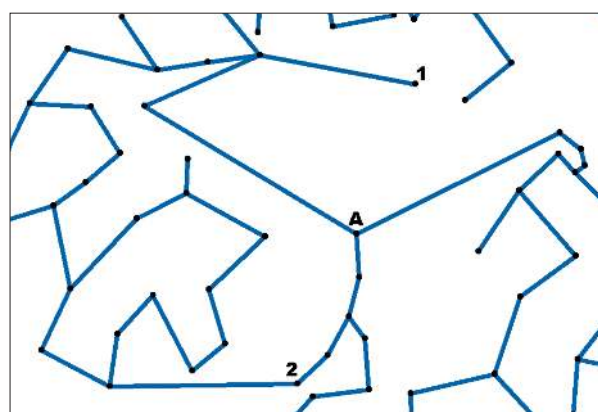


Figure 11. Sample display and similarity rating for *replication* trial of Experiment 2. Network distance is greater for A:1 than for A:2, while Node distance is greater for A:2. A:2 is rated as significantly more similar than A:1.

were also varied in 17 of the 65 trials that were essentially *replications* of the trials in Experiment 1. Because Network and Node distances were so strongly correlated in Experiment 1, we included these trials in order to more clearly establish whether Network distance rather than Node distance would be interpreted as indicating similarity relationships. The remaining 48 trials varied not only Network distances, but also a non-spatial visual variable. Of these trials, 15 varied the *width* of the blue network links.

Three levels of width were used: thin, medium, and thick (links in the *replication* trials were thick). Various trials displayed comparison links differing in width (all thick compared to all thin or all medium) or they were mixed in width (all thick compared to some mixture of the three widths) (Figure 8). Surrounding links (those not connecting the comparison documents) were mixed in width. In addition to the *replication* and *width* trials, another 15 trials kept the widths equal (thick) but varied the *value* of the blue network links. Three levels of value were used: light, medium, and dark (links in the *replication* trials were dark). Various trials displayed comparison links differing in value (all dark compared to all light or all medium) or mixed in value (all dark compared to some mixture of the three values) (e.g., Figure 9). Surrounding links were mixed in value. Finally, 18 trials kept the values and widths of the links equal (dark and thick) but varied the *hue* of the links. Blue and green links were used. Various trials displayed comparison links of the same hue (both all green), differing in hue (all blue compared to all green), or mixed in hue (all green compared to blue-green) (e.g., Figure 10). Surrounding links were mixed blue and green. The color schemes utilized for all displays in this experiment were selected using ©ColorBrewer, an online tool to help designers select perceptually

sound color schemes for maps and graphics (Brewer and Harrower 2002).

As in Experiment 1, participants first responded to five practice trials that did not involve distance comparisons. Participants also responded to 10 “warm-up” trials that introduced them to network displays varying in value, width, and hue. Participants answered 45 post-test questions after the main test trials, including the same 11 questions used in Experiment 1 about their personal backgrounds (however, they were administered before the main test trials in Experiment 1). The additional post-test questions were adapted from Experiment 1 to account for the new display types.

Procedure. Participants were tested as in Experiment 1. However, the 65 experimental trials in this experiment were not grouped into blocks; rather, they were presented in completely randomized orders to each participant.

Results and Discussion

Two participants, one female and one male, were excluded from the analyses because they answered ‘5’ (“A:1 and A:2 are equally similar”) to all trials of a given type (e.g., they rated all *replication* trials as ‘5’); this left a total of 33 participants. As in Experiment 1, similarity ratings were scored so that a mean rating less than 5 indicates that participants saw A:1 as more similar, while a mean rating greater than 5 indicates that participants saw A:2 as more similar.

Replication trials. Turning first to the *replication* trials, we find that the importance of Network distance over Node distance is clearly established. For example, Figure 11 depicts a trial in which Network distance is much greater (about three times) between A:1 than between A:2, but Node

distance is greater for A:2 than A:1 (three vs. two intervening nodes). A:2 were rated as significantly more similar than A:1 for this trial (mean = 6.6, sd = 2.3, $t[32] = 4.00$, $p < .001$). This shows that participants interpreted Network distance, and not Node distance, as relevant to similarity. Table 2 presents mean similarity ratings aggregated over trials that displayed the Network and Node distances A:1 and A:2 as equal or different.

The table shows that documents connected by links shorter in Network distance are rated as more similar, while documents connected by links equal in Network distance are rated as equally similar (not significantly different from 5.0). In contrast, documents connected by links equal in Node distance are rated as more similar if Network distances differ accordingly, while documents connected by links shorter in Node distance are rated as equally similar if Network distances do not differ. As further evidence of the importance of Network distance (rather than Node distance) in the replication trials, we examined the mean correlations of Network and Node distance with rated similarity for 17 trials. These correlations are moderate for Network distances: they are positive for 25 of the 33 participants, and average 0.38 across all participants. Based on a t-test calculated on the z-scores, this correlation is significantly greater than 0, $t(32) = 4.03$, $p < .001$. The correlation of rated similarity with Node distance is very weak and non-significant (-.07); the fact that it is negative again reflects that in some trials in which Node distance was equal, Network distance was unequal.

Width trials. In trials that varied link width, we found that this non-spatial visual variable greatly moderates the influence of Network distance. For example, Figure 12 depicts a trial in which Network distance is much greater (about three times) between A:2 than between A:1, but A:2 are connected by thick links while A:1 are connected by thin links. A:2 are rated as more similar than A:1 (mean = 5.9, sd = 2.3, $t[32] = 2.21$, $p < .05$). In these trials, participants did not equate Network distance with similarity; most interpreted thicker links as indicating greater similarity. Table 3 presents mean similarity ratings aggregated over trials that displayed the Network distances A:1 and A:2 as equal or different, and contrasted thicker links to thinner links or to links mixed in width.

The table shows that documents connected by links shorter in Network distance are not rated as being significantly more similar when the links differ in width, either homogeneously (e.g., all thick vs. all medium) or heterogeneously (e.g., all thick vs. mixed thick and medium). Documents connected

Type of Distance	Relationship of Distances A:1 and A:2	
	A:1=A:2	A:1<A:2
Network ^a	5.3	4.0***
Node ^b	5.2	6.1***

Notes. N=33 participants, total #trials=17. ^aA:1=A:2 for 5 trials, A:1<A:2 for 12 trials. ^bA:1=A:2 for 5 trials, A:1<A:2 for 12 trials. * $p < .05$, ** $p < .01$, *** $p < .001$, significance of difference from neutral similarity 5.0.

Table 2. Mean similarity ratings for replication trials in which network and node distances A:1 and A:2 equaled or differed, Experiment 2.




Figure 12. Sample display and similarity rating for *width* trial of Experiment 2. Network distance is greater for A:2 than for A:1, but A:2 is connected by thick links and A:1 by thin links. A:2 is rated as significantly more similar than A:1.

by links equal in Network distance but different in width are rated as being different in similarity; thicker links connote greater similarity than do thinner links, by over a full scale-point of similarity. The non-visual variable of link width essentially eliminated the relationship of Network distance to similarity judgments. Thicker was interpreted as more similar than thinner. As would be expected, the mean correlation of Network distance with rated similarity for the 15 *width* trials is close to 0 (.05) and non-significant. The correlations were positive for only 15 of the 33 participants.

Value trials. With regard to trials that varied link value, we found that this non-spatial visual variable also moderates the influence of Network distance, but not as strongly as does width. For example, Figure 13 depicts a trial in which Network distance is much greater (about three times) between A:1 than between A:2, but A:1 are connected by dark links while A:2 are connected by light links. The mean similarity rating of 5.4 (sd = 2.5) is not significantly greater than 5.0 for this trial. This shows that participants no longer

Width Relations of A:1, A:2	Relationship of Network Distances A:1 and A:2	
	A:1=A:2	A:1<A:2
A:1 thicker, A:2 thinner ^a	3.9***	5.2
A:1 thicker, A:2 mixed ^b	3.8***	5.5

Notes. N=33 participants, total #trials=15. ^a2 trials for =, 4 trials for <. ^b3 trials for =, 6 trials for <. *p < .05, **p < .01, ***p < .001, significance of difference from neutral similarity 5.0.

Table 3. Mean similarity ratings for width trials in which network distances A:1 and A:2 equaled or differed, Experiment 2.

simply equated Network distance with similarity; darker links were apparently interpreted by many participants as indicating greater similarity than did lighter links. Table 4 presents mean similarity ratings aggregated over trials that displayed the Network distances A:1 and A:2 as equal or different and contrasted darker links to lighter links or to links mixed in value.

The table shows that documents connected by links shorter in Network distance are *not* rated as being significantly more similar when the links differ in value, either homogeneously (e.g., all dark vs. all medium) or heterogeneously (e.g., all dark vs. mixed dark and medium). Documents connected by links equal in Network distance but different in value are rated as being a bit different in similarity; darker links connote greater similarity than do lighter or mixed links, though by only about a half scale point of similarity. This conclusion is bolstered by an examination of the mean correlation of Network distance with rated similarity for the 15 *value* trials. This correlation is quite small (.14) and non-significant. The correlations were positive for only 16 of the 33 participants. The non-visual variable of value weakened the effect of Network distance on similarity judgments, but unlike width, to a less consensual degree across participants. Darker tended to be interpreted as more similar than lighter, though several participants saw it the other way around.

Hue trials. Turning to the trials that varied link hue, we find that this non-spatial visual variable also moderates the influence of Network distance, but in a much less consistent way than does width or value. Figure 14 depicts four trials that demonstrate this. In (a), Network distance is greater for A:1 than for A:2, and both are connected by homogeneously green links. The fact that A:2 is rated as significantly more similar than A:1 (mean = 6.3, sd = 2.0, $t[32] = 3.83$, $p < .001$) suggests that Network distance is important to similarity judgments when hues are homogeneous, like the *replication* trials and the trials of Experiment 1.

Value Relations of A:1, A:2	Relationship of Network Distances A:1 and A:2	
	A:1=A:2	A:1<A:2
A:1 darker, A:2 lighter ^a	4.4*	4.9
A:1 darker, A:2 mixed ^b	4.6*	4.6

Notes. N=33 participants, total #trials=15. ^a2 trials for =, 4 trials for <. ^b3 trials for =, 6 trials for <. *p < .05, **p < .01, ***p < .001, significance of difference from neutral similarity 5.0.

Table 4. Mean similarity ratings for value trials in which network distances A:1 and A:2 equaled or differed, Experiment 2.

In (b), Network distance is again greater for A:1 than for A:2, but the connecting link structures are green for A:1 and blue for A:2. The fact that A:1 and A:2 are rated as equally similar (mean = 4.6, sd = 2.8, $t[32] = 0.82$, ns) suggests the moderating influence of hue on the relationship of Network distance to similarity; green is not preferred over blue, but the use of the two hues leads many participants to downplay the relevance of Network distance.

In (c), Network distance is equal for A:1 and A:2, but the connecting link structures are homogeneously green for A:1, and heterogeneously mixed green and blue for A:2. The fact that A:1 is rated as significantly more similar than A:2 (mean = 3.8, sd = 2.1, $t[32] = 3.25$, $p < .01$) suggests that homogeneous hues are interpreted as indicating greater similarity than are heterogeneously mixed links.

Finally, in (d), Network distance is greater for A:2 than for A:1, but A:2 are again connected by homogeneously green links while A:1 are connected by heterogeneous links that mix green and blue. The fact that A:1 is rated as significantly more similar than A:2 (mean = 4.1, sd = 2.4, $t[32] = 2.13$, $p < .05$) suggests the complex interrelationship of hue and distance to similarity; when Network distances differ, contrasts of hue homogeneity are ignored in favor of equating Network distance to similarity (although less strongly than when hues do not contrast).

Table 5 presents mean similarity ratings aggregated over trials that displayed the Network distances A:1 and A:2 as equal or different, and contrasted green links to green links, to blue links, or to links mixed in hue.

The table shows that documents connected by links shorter in Network distance are rated as significantly more similar when the links are the same hue. When links of one hue are compared to links of another hue, Network distance is ignored as a basis for similarity, but neither hue is interpreted consistently as reflecting greater similarity (at least when green is compared to blue). Finally, when the links differ heterogeneously (all green is compared to mixed green and blue), a more complex

pattern of similarity judgments results. Given equal Network distances, participants interpret homogeneous links as connecting more similar documents than do heterogeneous links. However, when Network distances are unequal, participants judge similarity rather inconsistently: Some base it on the distance relationships, others on the hue homogeneity. As a result, there is no consistent effect of Network distance or hue variations in such cases. Given the undeniable moderating influence of hue, albeit of a somewhat complex nature, it is not surprising that the mean correlation of Network distance with rated similarity for the 18 *hue* trials is small (.15) and non-significant. The correlations were positive for only 16 of the 33 participants.

Overall, participants in this experiment took a mean of 9.0 seconds (SD = 3.1 s) to respond to each network-display trial (the fastest mean response over all trials was 4.6 s, the slowest was 20.1 s). Responses differed somewhat as a function of the type of trials. They were fastest for the *replication* trials (8.5 s), slowest for the *hue* trials (9.5 s), and intermediate for the *width* and *value* trials (9.0 s and 9.1 s, respectively). Responses differed significantly in a repeated-measures MANOVA, $F(3, 30) = 6.14, p < .01$. This result corresponds well to the apparent complexity of similarity judgments participants had to make in the different trials, in particular the role the non-spatial visual variables played in moderating Network distance. Similarity in the *replication* trials depended only on Network distance. Similarity in the *width* and *value* trials depended on non-spatial visual variables in additive combinations with Network distance. Similarity in the *hue* trials depended on a complex combination of hue and Network distance.

Discussion

The results of the studies reported in this paper are consistent with findings from previous studies in that, all else being equal, participants indeed equate graphical inter-point distances with similarity in spatialized displays; in other words, we find empirical support for the operation of the distance-similarity metaphor in network spatializations (Fabrikant 2003; Montello et al. 2003). In addition, the outcomes of this research extend previous findings in several important ways. First, when viewing network spatializations, participants' default tendency is to equate metric distance along network links with similarity, rather than direct




Figure 13. Sample display and similarity rating for *value* trial of Experiment 2. Network distance is greater for A:1 than for A:2, but A:1 is connected by dark links and A:2 by light links. A:1 and A:2 are rated as equally similar.

distance across links or a topological measure of proximity such as the number of nodes.

Second, the magnitude of judged similarity is directly (negatively) correlated with network distance. That is, participants tend to linearly decrease their ratings of similarity between documents as the metric distance along the network between them increases. This finding provides support and a novel wrinkle to the hypothesized principle we called the "First Law of Cognitive Geography" (Montello et al. 2003), as expressed in the context of information visualization. By analogy to the well known First Law of Geography (Tobler 1970), the cognitive version states that people *believe* closer things to be more similar than distant things. For network spatializations, the First Law of Cognitive Geography apparently manifests along the links of the network rather than uniformly across 2D metric space.

Our finding that metric distance along links is the visual variable that denotes similarity in network displays has important design implications for network spatializations that are treated with graph layout algorithms, as well as for real-world transportation charts that are systematically distorted for legibility and efficiency reasons, such as many urban mass transit maps. If cognitive adequacy is one of the desired goals for designing network spatializations, then our findings suggest that link lengths should not be distorted purely for aesthetic reasons, as is the case with graph layout algorithms. People utilize the length of the links to assess similarity between information bearing items on a network spatialization. We may even speculate that if people explore maps of abstract spaces as if they explored maps of the real world, systematic distance judgment errors may




Figure 14. Sample display and similarity rating for *hue* trials of Experiment 2. The four panels depict the complex moderating effect of link hue on the relationship of network distance and similarity ratings.

occur in transportation charts that distort network distance for legibility. Underground stations that are closer on the network would be seen closer, even if the map designer did not intend network distances to accurately portray relationships due to layout reasons. Very little research (see Berendt et al. 1998) has tested people's spatial knowledge from studying transportation maps where link lengths have been purposefully distorted, and thus are not true to real-world network distance. Given the widespread use of such network maps in cities around the world, more empirical research on this topic is clearly needed.

Other major findings of our second experiment concern the effects of the non-distance visual variables of size (line width), color value (e.g., light blue vs. dark blue lines) and color hue (e.g., green vs. blue lines). These visual variables significantly moderate the distance-similarity metaphor in network spatializations. Our findings are gratifying in that they provide empirical support for certain cartographic design conventions for geographic maps that have been practiced for hundreds of years, and that were more recently explicitly formalized in the seminal writings of French cartographer Jacques Bertin

(1967; 1983); other cartographers have continued to update and revise these principles (Morrison 1974; MacEachren 1995).

Although often forgotten, the first of Bertin's seven visual variables is *location*—the location of the graphic mark with respect to the locations of other symbols on the map plane. In other words, the relative location of symbols in a graphic display is perceptually and thematically meaningful. This is in accordance with the Gestalt psychologists Wertheimer and Koffka (reviewed by Gregory (1987) and Goldstein (1989)) who posited that the arrangement of features in a picture or graphical image will influence the perceived thematic or group membership relations of elements. Our empirical results support these contentions in the context of information spatializations. The network emerges as a coherent feature from the background. Network structures become visually more salient (i.e., emerge as figure), which makes distances along them thematically more relevant for judging similarities than are direct distances across the open space between network structures (i.e., ground). In our work on point-display spatializations reviewed above (Montello et al. 2003), we also found an *emergent feature* effect:

when the background points against which people view comparison points is suitably non-uniform (i.e., has uneven groupings), people perceive linear and cluster features of points that modify the operation of the distance–similarity metaphor.

Visual variables are used in cartography to purposefully create figure-ground relationships, such that thematically relevant items become perceptually more salient. Bertin (1967; 1983) orders his remaining six variables according to the number of perceptual properties they carry. The visual variables of size and color value are the second and third most important variables in his list, and this is one reason we chose them in our research. Bertin contended that only symbols varying in size can be (approximately) perceived quantitatively. He thus claimed that size is the only appropriate visual variable to numerically depict ratios between signs according to their associated attribute magnitudes.

The direct perceptual association of size with magnitude is so strong, according to Bertin, that no recourse to a legend would be necessary to convey this. Furthermore, Imhof (1972) recommends that only line width should be used for quantitative network maps. Our findings do support the cartographic design principle that size is a powerful quantitative visual variable for designing cognitively adequate network spatializations. When available, and without recourse to a legend, participants used the visual variable of link width (e.g., size) to assess the magnitude of similarity between documents in network spatializations, essentially ignoring the visual variable of network distance.

When network links vary in width, the First Law of Cognitive Geography fails. It is notable, however, that the spatial metaphor of distance is supplanted by another spatial metaphor, in this case size. A variety of behavioral and cognitive scientists have discussed the cognitive and cross-cultural association of size and magnitude (Gattis 2001; Keating 1995; Lakoff 1987). For instance, the “larger/higher is more” metaphor is commonly found in everyday language, as in “the bank account is shrinking” or “the stock market is rising.” This spatial metaphor is one example of a set of image schemas that have been hypothesized to be at the core of human cognition (Lakoff, 1987).

Color value is the second dissociative visual variable in Bertin’s set. Dissociative variables cause the visibility of the sign on the visual plane to vary. For example, differently sized symbols take more or less white space on the visual plane, or different shades of gray cover more or less white space of the visual plane. Unlike size, Bertin (1967; 1983) claimed that value cannot be used to quantify variations in attribute

Hue Relations of A:1, A:2	Relationship of Network Distances A:1 and A:2	
	A:1=A:2	A:1<A:2
Both green ^a	5.5	3.6***
Green vs. blue ^b	4.8	5.0
Green vs. mixed ^c	3.8***	5.0

Notes. N=33 participants, total #trials=18. ^a1 trial for =, 2 trials for <. ^b1 trial for =, 2 trials for <. ^c4 trials for =, 8 trials for <. *p < .05, **p < .01, ***p < .001, significance of difference from neutral similarity 5.0.

Table 5. Mean similarity ratings for hue trials in which network distances A:1 and A:2 equaled or differed, Experiment 2.

magnitude. It is an ordered visual variable, that is, items can be logically ordered without recourse to a legend, but differences in magnitude cannot be measured visually, according to Bertin. We did not design our experiments to directly evaluate whether this is true. As a cartographer would expect, however, the variable of value did weaken the operation of the network distance–similarity metaphor, but not as much as line width did.

In essence, while most participants in our second experiment saw thicker to mean more similar, they were of less consensus in their opinion as to whether darker or lighter meant more similar. We interpret this finding to mean that even if value does not require a legend to establish ranks or intervals, a legend is still necessary to provide semantic context—to establish direction for the meaning of value. We contend that the “darker is more” or “lighter is more” principles are not uniformly self-evident in network spatializations; the maximum and the minimum values must be fixed in the legend to provide the direction (in a related vein, MacEachren and Mistrick 1992 found that whether a darker or lighter polygon is seen as figure or ground is rather inconsistent).

Although cartographers have often claimed that darker should be used for more (Brewer 1994; Dent 1999), we believe this should not be considered a consistent general principle. As an example of a strong exception, if one were designing a thematic map to show the magnitude of sunshine in a given area, it would surely be cognitively inadequate design to depict more days of sunshine by darkening a cartographic symbol. The ambiguity of semantic direction in most contexts, however, makes color value a weaker visual variable than size, and its moderation of the distance–similarity metaphor is in fact weaker.

The third non-distance visual variable we tested, color hue, is neither quantitative nor ordered, according to Bertin (1967; 1983). He considered it to be one of the strongest selective visual variables. Color may act as a simplifying or clarifying agent in map

design, affecting the perceptibility of a map (Dent 1999). Cartographers use color hue to denote associations in kinds of things (i.e., different groups or categories) for data collected at the nominal scale of measurement (e.g., land-use classes, language groups, soil types). Color has strong perceptual (harmonies, apparent advance or recede) as well as connotative effects (e.g., water is not red, fire is not blue) on the way maps are viewed and interpreted (Dent 1999). Not surprisingly for a cartographer, similar color hues strongly signal similar class membership, but the particular hue of a link connecting document points has no consistent moderating effect on the distance–similarity metaphor. This is because hue cannot be logically ordered or quantified visually. When available, participants used color hue to make similarity judgments, but due to the lack of a legend, they could not map particular hues onto similarity ratings in a consistent manner.

Design Implications for Network Spatializations

Elsewhere we have proposed a theoretical framework for cognitively adequate information spaces, grounded on GIScience as well as on cognitive, perceptual, and experiential principles. (Fabrikant and Skupin, in press). In this article, we elaborate on the importance of cartographic design principles (e.g., generalization, visual variables) for spatialization. Our current results provide empirical justification for this framework. According to MacEachren (1995), cartographic design theory has been subjected to only limited empirical verification.

As discussed above, graph layout algorithms such as Force-Directed Placement (FDP) are typically used to visualize Pathfinder Networks as shown in Figure 2 (e.g., Kamada and Kawai 1989). These algorithms are employed to identify node locations in a typically 2D Euclidean space by taking pair-wise proximity information as input. FDP algorithms optimize the visual output topologically by manipulating node placement, edge length, and node spacing based on quantifiable aesthetic criteria, for example, by minimizing edge crossings and striving for evenness of edge lengths so as to generate a visually pleasing layout.

In Figure 2 one can see how the edge length between some nodes of the force-directed graph contradicts the similarity strength computed between nodes by the Pathfinder Network Scaling algorithm (Schvaneveldt 1990), which is depicted using color value. The node placement is derived

from an FDP algorithm (i.e., Kamada and Kawai 1989) available in ©Interlink's KNOT software. The darker the blue edge color, the higher the similarity value according to the Pathfinder Network algorithm, but edge lengths produced by the FDP are not rendered accordingly. However, according to the distance–similarity metaphor, and confirmed by our experimental results, high similarity values should have shorter edges between nodes. If it is not possible to modify the algorithm, our results suggest that the spatialization designer could use line width to communicate similarity strength, or color value, provided there is a legend.

Conclusions and Prospects

Our empirical results suggest that the distance–similarity metaphor that has been found to work for point- and surface-display spatializations also applies to network spatializations, but in a different way. In the case of network displays, the distance–similarity metaphor operates by equating metric distance along network links to similarity. We find a negative correlation between network distance and similarity. The further away two points are on the network (in terms of distance, not node count), the less similar they are interpreted to be. Additionally, our results provide rare empirical evidence concerning how three visual variables that are widely used by cartographers on geographic maps, and generally in graphic design, are interpreted in network spatializations. We find that line size (width), value, and hue modify the distance–similarity metaphor in network-display spatializations in subtle yet logical ways. Our empirical findings thus extend initial findings for the application of sound cartographic design principles for point and surface displays to network-display spatializations.

By providing design guidelines to information visualization developers based on empirical data and theories about people's similarity perceptions in network displays and other types of spatializations, we can ensure that cognitively adequate 2D and 3D information spatializations can be constructed.

ACKNOWLEDGMENTS

Generous funding was provided by the National Geospatial-Intelligence Agency (NGA). We thank Amy Linker and Ben Turner for their assistance with data collection, and the UCSB students who were willing to participate in our research. We also thank David Mark and Corrine Jörgensen at SUNY Buffalo for their collaboration on this project since its inception.

REFERENCES

- Berendt, B., R. Rauh, and T. Barkowsky. 1998. Spatial thinking with geographic maps: An empirical study. In: H. Czap, H. P. Ohly, and S. Pribbenow (eds), *Herausforderungen an die Wissensorganisation: Visualisierung, Multimediale Dokumente, Internetstrukturen*, Ergon, Würzburg, Germany: 63-73.
- Bertin, J. 1967. *Sémiologie graphique: les diagrammes – les réseaux – les cartes*. Mouton, Paris.
- Bertin, J. 1983. *Semiology of graphics*. University of Wisconsin Press, Madison, Wisconsin.
- Börner, K., C. Chen, and K.W. Boyack. 2003. Visualizing knowledge domains. *Annual Review of Information Science and Technology*, vol. 37, no. 178-255.
- Brewer, C. A. 1994. Color use guidelines for mapping and visualization. In: A. M. MacEachren, and D.R. Fraser Taylor (eds), *Visualization in modern cartography*. New York, New York: Elsevier. pp. 269-78.
- Brewer, C. A., and M. Harrower. 2002. ColorBrewer. A web tool for selecting colors for maps: ©, Cynthia Brewer, Mark Harrower, and The Pennsylvania State University. All rights reserved.
- Card, S. K., J. D. Mackinlay, and B. Shneiderman. 1999. *Readings in information visualization. Using vision to think*. San Francisco, California: Morgan Kaufmann.
- Chen, C. 1999. *Information visualisation and virtual environments*. London, U.K.: Springer.
- Chen, C. 2003. *Mapping scientific frontiers: The quest for knowledge visualization*. London, U.K.: Springer.
- Deerwester, S., S. T. Dumais, G.W. Furnas, T.K. Landauer, and R. Harschman. 1990. Indexing by latent semantic analysis. *Journal of the American Society of Information Science* 41(6): 391-407.
- Dent, B. D. 1999. *Cartography: Thematic map design*. Dubuque, Iowa: Wm. C. Brown,
- Fabrikant, S. I. 2001. Visualizing region and scale in semantic spaces. In: *Proceedings, The 20th International Cartographic Conference, ICC 2001*, Beijing, China, Aug. 6-10, 2001. pp. 2522-9.
- Fabrikant, S. I. 2003. Distanz als Raummetapher für die Informationsvisualisierung. Distance as a spatial metaphor for the visualisation of information. *Kartographische Nachrichten* 52(6): 276-82.
- Fabrikant, S. I., and B.P. Buttenfield. 2001. Formalizing semantic spaces for information access. *Annals of the Association of American Geographers* 91(2): 263-80.
- Fabrikant, S. I., and A. Skupin. Cognitively plausible information visualization. In: J. Dykes, A. M. MacEachren, and M.J. Kraak (eds), *Exploring GeoVisualization*. Amsterdam, The Netherlands: Elsevier. (In press).
- Garland, K. 1995. *Mr. Beck's Underground map*. Capital Transport, Middlesex, U.K.
- Gattis, M. 2001. *Spatial schemas and abstract thought*. Cambridge, Massachusetts: The MIT Press.
- Goldstein, E. B. 1989. *Sensation and perception, 3rd ed.*, Belmont, California: Wadsworth,
- Gregory, R. L. 1987. *The Oxford companion to the mind*. Oxford. U.K.: Oxford University Press.
- Howell, D. C. 2002. *Statistical methods for psychology*, 5th ed. Pacific Grove, California: Duxbury/Thomson Learning.
- Imhof, E. 1972. *Thematische Kartographie*. Berlin/New York: DeGruyter.
- Kamada, T., and S. Kawai. 1989. An algorithm for drawing general undirected graphs. *Information Processing Letters* 31(1): 7-15.
- Keating, E. 1995. Spatial conceptualizations of social hierarchy in Pohnpei, Micronesia. In: A. U. Frank, and W. Kuhn (eds), *Spatial information theory: A theoretical basis for GIS*. Lecture Notes in Computer Science 988. Berlin, Germany: Springer Verlag. pp. 463-74.
- Lakoff, G. 1987. *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago, Illinois: University of Chicago Press
- MacEachren, A. M. 1995. *How maps work. Representation, visualization, and design*, New York, New York: Guilford Press.
- MacEachren, A. M., and T. A. Mistrick. 1992. The role of brightness differences in figure-ground: Is darker figure? *The Cartographic Journal* 29: 91-100.
- Montello, D. R. 1991. Spatial orientation and the angularity of urban routes: A field study. *Environment and Behavior* 23: 47-69.
- Montello, D. R., S. I. Fabrikant, M. Ruocco, and R. Middleton. 2003. Spatialization: Testing the First Law of Cognitive Geography on point-spatialization displays. In: W. Kuhn, M.F. Worboys, and S. Timpf (eds), *Spatial information theory: Foundations of geographic information science. Conference on Spatial Information Theory (COSIT) 2003*. Lecture Notes in Computer Science 2205. Berlin, Germany: Springer Verlag. pp. 335-51.
- Morrison, J. L. 1974. Changing philosophical-technical aspects of thematic cartography. *The American Cartographer* 1(1): 5-14.
- Ruggles, A. J., and M. P. Armstrong. 1997. Toward a conceptual framework for the cartographic visualization of network information. *Cartographica* 34: 33-48.
- Schvaneveldt, R. W. (ed.). 1990. *Pathfinder associative networks: Studies in knowledge organization*. Norwood, New Jersey: Ablex.
- Tobler, W. R. 1970. A computer movie simulating urban growth in the Detroit region. *Economic Geography* 46(2): 234-40.
- Tversky, B. 1981. Distortion in memory for maps. *Cognitive Psychology* 13: 407-33.
- Tversky, B. 2000. Some ways that maps and diagrams communicate. In: C. Freksa, W. Brauer, C. Habel, and K.F. Wender (eds), *Spatial Cognition II, Integrating abstract theories: Empirical studies, formal methods, and practical applications*. Lecture Notes in Computer Science 1849. Berlin, Germany: Springer Verlag. pp. 72-9.
- Wise, T. A. 1999. The ecological approach to text visualization. *Journal of the American Society of Information Science* 53(13): 1224-33.