

Visualizing space, time, and agents: production, performance, and preference

Angela Kessell · Barbara Tversky

Received: 13 June 2010 / Accepted: 22 October 2010 / Published online: 17 November 2010
© Marta Olivetti Belardinelli and Springer-Verlag 2010

Abstract Visualizations of space, time, and agents (or objects) are ubiquitous in science, business, and everyday life, from weather maps to scheduling meetings. Effective communications, including visual ones, emerge from use in the field, but no conventional visualization form has yet emerged for this confluence of information. The real-world spiral of production, comprehension, and use that fine-tunes communications can be accelerated in the laboratory. Here, we do so in search of effective visualizations of space, time, and agents. Users' production, preference, and performance aligned to favor matrix representations with time as rows or columns and space and agents as entries. Overall, performance and preference were greater for matrices with discrete dots representing cell entries than for matrices with lines, but lines connecting cells may provide an advantage when evaluating temporal sequence. Both the diagram type and the technique have broader applications.

Keywords Diagram · Production · Comprehension · Preference · Space · Time

This research was carried out as part of the first author's doctoral dissertation.

A. Kessell
Stanford University, Stanford, CA 94305, USA

Present Address:
A. Kessell (✉)
San Jose State University, NASA Ames Research Center,
Mail Stop 262-4, Moffett Field, San Jose, CA 94035, USA
e-mail: angela.kessell@nasa.gov

B. Tversky
Columbia Teachers College, New York, NY, USA

Introduction

Arranging, understanding, or tracking people or objects in space and time are ubiquitous problems in science, business, and everyday life. Geographers, anthropologists, and historians study the migrations of people over time; botanists, the flow of pollen and plants; and epidemiologists, the spread of diseases. Scheduling, of social events, conferences, or deliveries, requires organizing people or goods in places and times. Not just history and planning, but also our memories are organized around people, places, and time (e.g., Taylor and Tversky 1997; Wagenaar 1986; Zwaan and Radvansky 1998). The quantity of information to organize and track can quickly grow. As for many complex problems, visualizations can help. Visualizations can group and simplify information and array it spatially in natural ways to transfer human agility in making spatial comparisons and spatial inferences to making abstract comparisons and inferences (e.g., Larkin and Simon 1987; Tversky 2001).

Although there are common visualizations of some aspects of space, time, and people or objects, many are not suitable for organizing the information precisely. Weather maps, for example, show general movements of storms or pressure systems over space. The typical maps showing population or disease migration are similar. Such displays map space geographically, a natural mapping that is readily comprehended, but they do not afford exact spatial-temporal resolution. Graphs of, for example, changes in population or GNP over time for varying geographical regions come closer. They lose the two-dimensional spatial correspondences of cartographic maps, but they provide spatial-temporal resolution of changes in "object," in this example, population or GNP. Such displays use lines to connect the temporal values of population or GNP. Lines not only draw

the eye and connect the values, but also encourage thinking in terms of trends (Zacks and Tversky 1999). However, by drawing the eye to, and thus emphasizing, one variable, they may inhibit inferences on the others. Of course, in many contexts, this is precisely the advantage visualizations offer over other kinds of information displays, deemphasizing context-irrelevant information and highlighting key information for a particular audience. Still, many cases do not demand continuous temporal or spatial resolution. Yearly GNP for an entire country or daily weather for an entire region may be sufficient. Still other cases involving location, time, and people/objects, such as scheduling a series of meetings, are inherently categorical. A canonical visualization for this set of variables has not evolved.

From a broader perspective, space, time, and people/objects could be viewed abstractly, as three data sets, data that may be a mixture of interval, ordinal, and categorical. However, the human mind does not treat space, time, people, and objects as abstract entities devoid of content. On the contrary, there are deeply rooted cognitive expectations for how these particular variables should be spatially arrayed (e.g., Tversky et al. 1991). Clues to effective visualizations can come from the natural cognitive correspondences revealed in people's spontaneous graphic productions (e.g., Heiser et al. 2004; Hurley and Novick 2006; Novick et al. 1999; Tversky et al. 2000, 2006). Although the spatial cognitive correspondences for some of these variables have been uncovered, there has been no research to uncover their joint cognitive correspondences. For example, temporal relations are typically mapped horizontally, both in language and in diagrams, corresponding to writing direction (e.g., Boroditsky 2001; Clark 1973; Tversky et al. 1991). Space is typically mapped both horizontally and vertically, corresponding to the perceived two-dimensional space of the flattened world (Tversky 1981). How, then, would people map both space and time, and people in addition, in a single two-dimensional display?

Effective visualizations, such as sketch maps, like most effective communication devices, have developed in a community of users through a kind of informal user-testing. Some people create visualizations, and others try to understand them or use them. If the visualizations do not work well, they are fine-tuned, an iterative process similar to collaboration in conversation (Clark 1996). This spiral process, of production, comprehension, reproduction, re-comprehension, can be brought into the laboratory and accelerated. In essence, the laboratory technique turns users into designers by employing three key tasks: production, preference, and performance. Here, we apply the *user-as-designer method* in the service of effective visualizations of space, time, and people/objects. In the previous *user-as-designer* cases, route maps and assembly instructions, the

community of users had had extensive experience with the tasks, and the three measures converged (Tversky et al. 2006; Heiser et al. 2004).

We chose a task that is representative of those requiring organization of space, time, and people/objects that would be familiar and accessible to our participants, college students. Their task was to keep track of the changing locations of students at different times of the day, a task, like wayfinding and assembly, in which they could be regarded as experts. To keep the task manageable, people and place were categorical and time ordinal, and there were four of each students, places, and times of the day. Information about people, places, and times is more like a map than like a route, that is, the information in such displays can be used for a broad range of inferences, such as individual habits, collocations of individuals, popularity of places over time, and more.

Study 1: production of visualizations

Methods

Participants

One hundred and seventy-two undergraduates from Stanford University participated in the study in return for partial course credit.

Stimuli

Participants completed a questionnaire with the following instructions:

Suppose you need to keep track of where your friends are during the day so that you can answer questions like: Were Alex and David together downtown? Or, how many people were at the gym in the afternoon? Did Justin go to the library before he went to the movies?

Here is the information you will have to keep track of:

People: Alex, David, Justin, Sammy

Places: dorm, library, downtown, gym

Times: morning, noon, afternoon, night

Using the space below, please invent a way to display that information so you can answer those kinds of questions as quickly and accurately as possible.

The presentation order of the three variables, people, places, and times, was counterbalanced across participants.

Participants worked on the questionnaire for as long as they liked, which was typically around 5 min.

Results

Twelve participants failed to follow instructions and were excluded from further analyses, leaving 160 usable responses. Presentation order of people, places, and times had no effects ($P_s > .10$), so the results were combined.

Two coders coded the visualizations as matrices, lists, textual narratives, and other, with inter-rater agreement of $\rho = .86$ ($P < .001$). The vast majority, 75%, were matrices. Most were simple matrices with each variable appearing once, but about a quarter were coded as extended, nested, or multiple, in which variables appeared more than once. Fifteen percent were hierarchical lists and trees, 6% were narratives, and the remaining four percent could only be classified as “other.”

The simple matrices had a fairly stereotypic organization, as evident from Fig. 1. Time almost always appeared as an axis ($\chi^2_2 = 24.21$, $P < .001$), though equally frequently on the x - and y -axes ($\chi^2_1 = 1.03$, $P > .10$). People and places were equally likely to appear in the cells ($\chi^2_1 = 0.11$, $P > .10$) and equally likely to appear on the x -axis ($\chi^2_1 = 0.21$, $P > .10$).

Discussion

Matrices, or tables, were by far the most commonly produced visualization for keeping track of people in various places at different times. Matrices are apt representations for this situation because the data are categorical, and matrices allow all values of one variable to be combined with all values of another variable, for a factorial combination of the values across two variables (Hurley and Novick 2006; Novick 2006). Lists and trees are not as

appropriate because they only represent positive information, that is, the occurrence of a person at a place at a time. Lists and trees, although structured, have no inherent length or branching constraints, and so it may not be obvious whether further content should be added or where it should go in the diagram. Matrices, on the other hand, highlight negative as well as positive information, by empty cells. Empty cells are visibly apparent and call attention to an eventless time or place, as well as to the possibility of missing information. Furthermore, because of their structure, lists and trees can imply an unintended hierarchy among the variables. Narratives fail to take advantage of spatial organization to organize the information.

There was a strong tendency to put time along one of the axes rather than in the cells. This may be due to the fact that time was the only ordinal variable, and this temporal order likely suggested a linear spatial organization (Clark 1973; Tversky et al. 1991). The other two variables, people and places, were categorical and therefore more conceptually congruent to cells than axes.

The omissions are also noteworthy. No one produced a line graph and no one produced a map-like display, although the locations were known to participants. One reason for the latter is that maps use two dimensions, making it difficult to also represent people and time. Another reason is that the spatial relationship of the particular locations named was not salient or meaningful in the present case. The categorical and cross-categorical nature of the variables seems to encourage a matrix visualization.

Study 2: queries: performance and preference

Visualizations of movements of agents in space and time should allow answering a broad range of queries, where people were at certain times, what places are popular at what times, collocations of certain individuals, temporal sequence of visiting certain locations, and more. In the next experiment, we turn to performance, namely retrieval time, to find answers to a range of queries. In the production study, participants showed a clear preference for matrices, with time along one of the axes, for organizing and keeping track of information about people, places, and times. But what kinds of information retrieval tasks do these space-time-agent matrices support? Although participants favored matrices, their productions did not converge to favor other aspects of the visualization. One issue is the entries of the matrices. Participants used a broad range, including X's, check marks, initials, and more. Here, we compared entries of color-coded, filled circles to another graphic device common to other displays of time, lines, again color-coded, and connecting the elements in the cells. Whereas

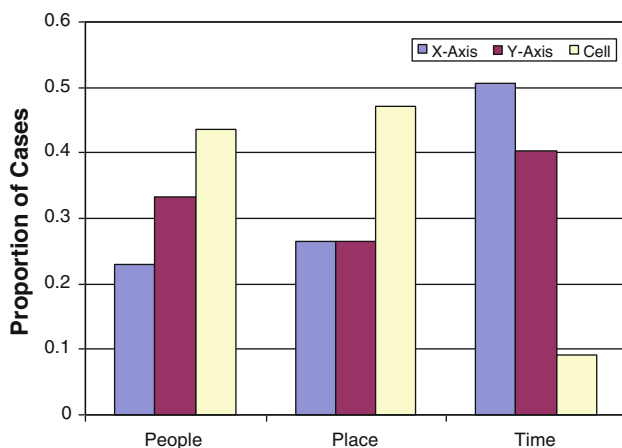


Fig. 1 Proportion of cases in which people, place, and time were displayed in x -axis, y -axis, and cell position in simple matrices

information on the axes (rows and columns) is laid out linearly, making it easy to organize and track, the information in the cells is not, so connecting lines might help guide attention and facilitate organization and tracking.

How might these two displays, one using filled circles and the other using lines, affect performance and preference? As discrete elements, filled circles, or dots, should facilitate discrete comparisons and verifications of quantity, such as “greater than” and “three people”. In contrast, lines form paths, connect entities, and show intersections. Lines should therefore be useful for spotting trends, tracking entities over time, and identifying intersecting paths (Zacks and Tversky 1999).

A second variable of interest here was the arrangement of time, place, and agent along axes and in cells. The productions placed time along an axis but did not definitively favor horizontal over vertical. This contrasts to previous work that showed a preference for time, a neutral dimension, displayed horizontally (Tversky et al. 1991). Here, we also test performance and preference for horizontal and vertical time.

Performance methods

Participants

One hundred and twenty-eight undergraduates from Stanford University participated in the study in return for partial course credit or for pay.

Stimuli

The stimuli were matrices displaying information about the locations of four men at four different times of the day. As in the production study, the variables were categorical (people and places) and ordinal (times), not interval or scale. Below the matrix was a query. The task of participants was to decide, as quickly and accurately as possible, whether the statement of the query was true or false based on the information in the matrix (Fig. 2).

	Morning	Noon	Afternoon	Evening
Dorm	●		● ●	
Library	●	●		
Bookstore	●	●		● ● ●
Gym	●	● ●	●	

Alex
 David
 Justin
 Sammy

David went to the gym some time after Sammy went to the gym.

Fig. 2 Screenshot of one experimental trial

Three properties of the matrices, the axes, elements, and schedule, were manipulated in a $2 \times 2 \times 2$ mixed design. For half of the participants, time appeared on the horizontal axis, and for the other half, time appeared on the vertical axis (Fig. 3). For all participants, half of the trials used dots to represent people and the other half used lines. The dots and lines appeared in four different colors, and a key located to the right of the matrix connected each color with a different person (Fig. 2). Also for all participants, half of the trials used one schedule, or assignment of the four people to the four places over the four time periods, and the other half used another schedule (Fig. 3). Each schedule employed a different color-person mapping, and the schedule-element pairing was counterbalanced across participants.

The experiment was programmed in E-Prime 1.0 (Psychology Software Tools, Pittsburgh, PA), and stimuli were presented on a 19-inch, Dell Ultrasharp flat panel monitor.

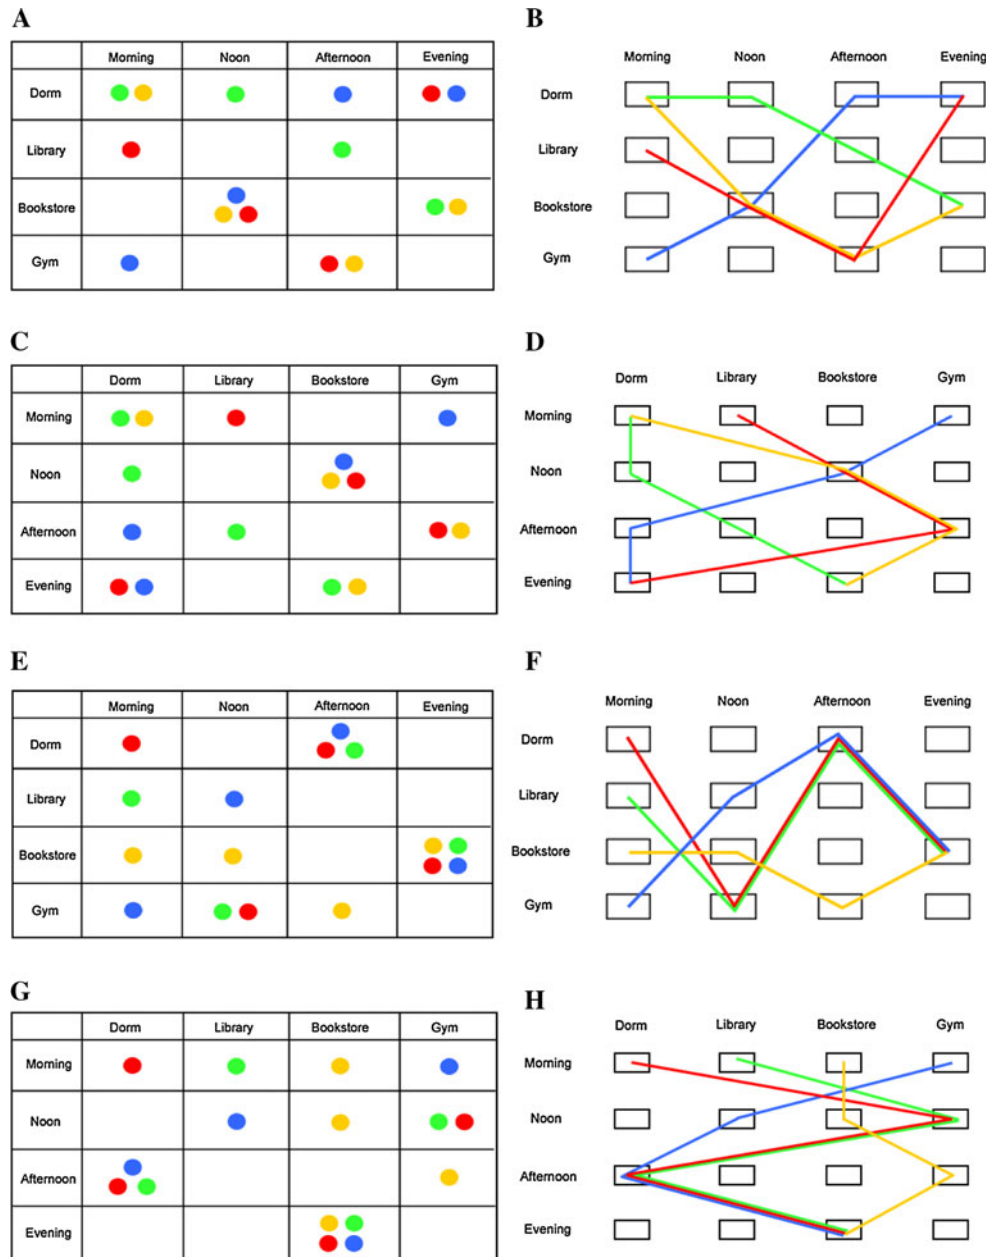
Design

The experiment was broken into two blocks, each consisting of four practice trials and 52 experimental trials. For one entire block, a single matrix, with a particular element type, axis assignment, and schedule, was displayed and participants verified statements, one at a time, based on the information contained in that matrix. For the second block, a different matrix was displayed, this time using the other element type and schedule, and all statements were verified according to the information in this new matrix. The axis assignment remained the same across blocks and only varied across participants. The practice trials used two additional schedules, not used in the experimental trials. Participants were instructed to respond as quickly yet accurately as possible. Participants pressed one key if the statement was true and another if it was false. Response time and accuracy were recorded by E-Prime.

Queries

There were two sets of 52 query statements, one for each experimental schedule, and eight different types of statements, each with a different focus: Person; Place; Time; Place and Time; People, Place, and Time; Sequence; Intersection; and Miscellaneous (a catch-all category where we explored a few additional statement types). The category membership rules and sentence order for each query type are listed in Table 1. Within each category, half of the query statements were false. An attempt was made, with partial success, to equate the number of syllables between and especially within categories. Across the two schedules, the wording of the query statements was kept as similar as possible, while still maintaining a .5 ratio of true statements. Table 1 also presents the predicted patterns of

Fig. 3 Experimental stimuli: **a, b** Schedule 1, time on *horizontal* axis; **c, d** Schedule 1, time on *vertical* axis; **e, f** Schedule 2, time *horizontal*; **g, h** Schedule 2, time *vertical*



response times for each query type. The predictions are based on the hypotheses that time should be processed more efficiently when it is arrayed horizontally and that dots should facilitate verification of people and places, but that lines connecting agents should facilitate temporal comparisons.

Performance results

Overall response time

All analyses of response time include only correct trials. Overall mean response time was 6.6 s with a standard

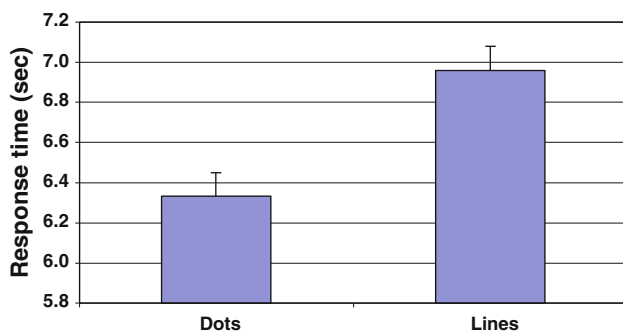
deviation of 1.2 s. A repeated measures analysis of variance (ANOVA) with within-subject variable element and between-subject variable axes revealed no element by axes interaction effect and no main effect of axes ($F_s < 2.41$; $ps > .10$; Fig. 4). There was a main effect of element ($F(1,126) = 52.29, P < .001$; Fig. 4) such that dots led to faster response times overall.

Response time by query type

A repeated measures ANOVA of response time with within-subject variable element and between-subject variable axes was performed for each query category for which we had an

Table 1 Number of statements, membership criteria, sentence order, example statement(s), and response time predictions for each query statement category

Category	<i>N</i>		
Person	12	Rule	Always and only name a specific person.
		Order	Always start with a person's name.
		Example	Sammy was with someone at all times.
		Predictions	Dots = Lines; Time <i>x</i> -axis < Time <i>y</i> -axis
Place	16	Rule	Always and only name a specific place.
		Order	Half start with place, half with person or time.
		Examples	The dorm was always occupied. All four people visited the bookstore.
		Predictions	Dots < Lines; Time <i>x</i> -axis < Time <i>y</i> -axis
Time	16	Rule	Always and only name a specific time.
		Order	Half start with time, half with person or place.
		Examples	In the morning, people were in four places. Three people were together at noon.
		Predictions	Dots < Lines; Time <i>x</i> -axis < Time <i>y</i> -axis
Place and Time	16	Rule	Always and only name a specific time and place.
		Order	Half start with place, half with time.
		Examples	The library had two people in the evening. In the morning, one person was at the gym.
		Predictions	Dots < Lines; Time <i>x</i> -axis < Time <i>y</i> -axis
People, Place, and Time	8	Rule	Always name a specific person, place, and time.
		Order	Half start with people, half with time.
		Examples	David was in the dorm in the evening. At noon, Alex was in the dorm.
		Predictions	Dots < Lines; Time <i>x</i> -axis < Time <i>y</i> -axis
Sequence (person and multiple places)	12	Rule	Always and only name a specific person and places.
		Order	Always start with a person's name.
		Example	Alex went to the dorm sometime after the library.
		Predictions	Lines < Dots; Time <i>x</i> -axis < Time <i>y</i> -axis
Intersection (multiple people)	12	Rule	Always and only name at least two specific people.
		Order	Always start with a person's name.
		Example	Alex and Sammy met more than once.
		Predictions	Lines < Dots; Time <i>x</i> -axis < Time <i>y</i> -axis
Misc.	12	Examples	Sammy and David both avoided the library. Alex visited all the places Justin visited.

**Fig. 4** Response times by graphic element type. Error bars represent standard error of the mean

a priori prediction. For all query categories other than Sequence (Fig. 5), overall response times were shorter with dots than with lines (Person $F(1,126) = 26.04$, $P < .001$; Place $F(1,126) = 66.66$, $P < .001$; Time $F(1,126) = 28.41$, $P < .001$; Place and Time $F(1,126) = 12.46$, $P < .01$; Intersection $F(1,126) = 28.04$, $P < .001$; People, Place, and Time $F(1,126) = 4.15$, $P < .05$). These effects were qualified by an element by axes interaction for both Time and People, Place, and Time query statements. Specifically, when time appeared on the horizontal axis, the advantage for dots was only marginal for Time query statements ($t(63) = 1.72$, $P < .10$) and disappeared




Fig. 5 Response times by graphic element type for Sequence statements. Error bars represent standard errors of the means

altogether for People, Place, and Time statements ($t(17) = 0.2, P > .10$).

Looking just at the effects of axis assignment, for Person query statements and Time query statements, response times were marginally shorter when time was on the horizontal axis (Person $F(1,126) = 3.39, P < .10$; Time $F(1,126) = 4.45, P < .10$). In contrast, for Place query statements, response times were marginally shorter when time was on the vertical axis ($F(1,126) = 3.54, P < .10$).

Overall accuracy

Overall mean accuracy was 0.94 with a standard deviation of 0.03. A repeated measures ANOVA of accuracy with within-subject variable element and between-subject variable axes revealed no element by axes interaction effect and no main effect of axes ($F_s < 1.17; p_s > .10$). There was a main effect of element ($F(1,126) = 4.73, P < .05$) such that dots led to higher accuracy overall.

Accuracy by query category

Again, repeated measures ANOVAs were performed. The only difference found was that for Place and Time query statements responses were marginally more accurate with lines than with dots, possibly the result of a speed-accuracy trade-off ($F(1,126) = 3.38, P < .10$).

Performance discussion

In retrieving information in response to queries about people, places, time, and various combinations, people responded faster and more accurately using matrices with dots than matrices with lines with one important exception. The advantage of dots over lines in information retrieval disappeared when participants verified queries about temporal sequences, a result in line with predictions. The lack of differences in accuracy is likely due to ceiling effects, as overall accuracy was high. Although participants were encouraged to respond as quickly as possible, they had unlimited time to respond to each query statement.

Restricting retrieval time might have made accuracy a more sensitive measure in this experiment.

Preference methods

Participants

Immediately after participating in the performance task, all 128 participants completed the preference task.

Stimuli and design

Participants filled out a multiple page questionnaire (Fig. 6). At the top of each page were two color matrices, one with lines and one with dots. The assignment of the axes matched the condition participants saw in the performance task. The first page contained the following instructions:

Today you used two different visual representations of the schedules of four people. These are labeled A and B below. We are interested in your thoughts on the relative advantages and disadvantages of each. Below are some sample statements like the ones you saw in the experiment. For each statement, please indicate whether you found one representation to be more useful than another in deciding whether this statement was true or false. Also, please feel free to explain why.

There were eleven statements in total, with one to two examples from every category except Miscellaneous. Participants responded by circling “A,” “B,” or “No preference.”

Preference results

Overall, there were no differences in preferences by axis assignment ($\chi^2 = 1.13, P > .05$). For nearly all statements, participants preferred the dots over the lines ($\chi^2_s > 34.33; p_s < .001$; Fig. 7). However, participants preferred the lines over the dots for the two Sequence statements ($\chi^2_s > 122.14; p_s < .001$; Fig. 8).

Preference discussion

Overall, participants’ preferences matched their performance and matched the previous participants’ productions. Participants preferred dots over lines in general, but preferred lines over dots for evaluating Sequence statements.

Discussion and conclusions

Visual communications are some of the oldest as well as the newest form of communication, from local maps

Fig. 6 Partial page from the preferences questionnaire



4. The dorm was always occupied.

Preferred: A B No preference

Reason(s): _____

5. At noon, people were in more than one place.

Preferred: A B No preference

Reason(s): _____




Fig. 7 Participants' graphical element preferences for all statements types but Sequence statements




Fig. 8 Participants' graphical element preferences for the two Sequence statements, "Justin went to the dorm sometime after the library," and "David went directly from the library to the dorm to the bookstore."

inscribed in stone to glitzy graphs in daily papers. They have been found in diverse cultures and preceded written language. Despite their prevalence, visualizations of the

non-visible, notably data, are a recent phenomenon. Although many, such as the periodic table or Minard's depiction of Napoleon's unsuccessful campaign against Russia or Beck's London tube map, are praised for their clarity, others are opaque and confusing. Designing good graphics is a challenge. Many common graphics, such as route maps, developed through practice in a community of users, a process that fine-tunes the graphics and improves communicative efficacy. That process can be accelerated in the laboratory and used to reveal design principles in specific domains (Heiser et al. 2004). Users produce visualizations, and those visualizations are tested for performance and preference in other users. When there is convergence, that is, when the same visualizations prevail in all three tasks, design principles can be extracted.

This user-as-designer method was applied here to a common visualization problem, simultaneously communicating information about space, time, and agent, using the paradigmatic task of visualizing the changing locations of agents in time. The production experiment yielded strong consensus on a tabular or matrix representation of that information. Moreover, respondents put time on an axis but placed people and locations equally on an axis or in the cells.

The performance and preference tasks assessed two main variants and two subvariants of matrices. The performance task entailed verification of many kinds of

queries that could be made from the visualizations, the kinds of information that might naturally be sought. Were certain people in the same place at the same time? Did a person go from one specific place to another? Did people congregate in a particular place? In the main variant of the visualization, table entries were color-coded dots; in the other, they were color-coded lines connecting cell entries. In both cases, cell entries were people. The visualizations also varied on the orientation of time, vertical or horizontal. To assess performance, participants verified whether a large set of relations among space, time, and agent were true or not, using the different variants of the visualization. Afterward, they were asked which visualization they preferred for verifying a range of queries of space, time, agent, and combinations.

Because they connect, lines were expected to facilitate one important kind of query about people, their temporal paths among places. This prediction derived from previous work where lines were used to connect places or values and encouraged trend interpretations (Tversky and Lee 1998, 1999; Zacks and Tversky 1999). The lines connected the people across time and place. Lines, however, clutter, making it more difficult to discern cell entries. Hence, lines should interfere with other comparisons; they should be facilitated by unconnected dots as cell entries. The predictions for the orientation of time are less clear. On the one hand, earlier work on spontaneous graphic productions in children and adults found that they arrayed time horizontally more frequently than vertically. This is, of course, the standard in graphs. Moreover, horizontal is often preferred for neutral dimensions, such as time, and the vertical for evaluative dimensions, such as preference, perhaps because the horizontal dimension is more neutral than the vertical. However, a minority of participants in that study did map time from up/early to down/late, as in calendars and date books (Tversky et al. 1991) or some temporal expressions in Mandarin Chinese (Boroditsky 2001).

Performance was better for matrices with dots than for matrices with lines for all but one of the statement types. The exception was queries of temporal sequence. As for the orientation of time, verification performance was equally good when time was vertical as when it was horizontal. Participants' preferences were in alignment with their performance. That is, they preferred matrices with dots for all question types except for temporal sequence, and these preferences held whether time was oriented vertically or horizontally.

Thus, production, performance, and preference for visualizing time, space, and agents converged on a tabular representation with dots for entries, except for temporal sequences, where lines were as good and may be better. There was also convergence that time should be an axis,

though there was no preference for the orientation of the axis, vertical or horizontal. This may be because there are strong correspondences of time to both vertical and horizontal in common visualizations. Graphs typically plot time horizontally, as do time lines. However, calendars and date books display time vertically, with earlier times at the top and later ones at the bottom. Because participants have undoubtedly used both tools, they are used to imagining time both vertically and horizontally.

The users-as-designers paradigm has produced consensus on visualizations of time, space, and agent. This paradigm has been used successfully in other domains, namely diagrams of routes and assembly (Heiser et al. 2004). Together, the projects provide support for this set of procedures as a general tool for designing tools for human use. There are undoubtedly limitations to this design tool. In all of these situations, users had some familiarity with the task and with visualizations of the task. Thus, there is reason to believe that users can serve as effective designers when they have some domain expertise, that is, past experience in comprehending and perhaps even in producing relevant visualizations and in thinking more abstractly about the information. All things equal, co-opting users to be designers and evaluating their productions, performance, and preferences looks promising as part of a program for creating effective designs.

Acknowledgments We are grateful for the support of the Stanford Regional Visualization and Analysis Center, NSF REC-0440103, NSF IIS-0725223, NSF IIS-0855995, and NSF HHC 0905417.

References

- Boroditsky L (2001) Does language affect thought? Mandarin and English speakers' conceptions of time. *Cogn Psychol* 43:1–22
- Clark HH (1973) Space, time, semantics, and the child. In: Moore TE (ed) *Cognitive development and the acquisition of language*. Academic Press, New York
- Clark HH (1996) *Using language*. Cambridge University Press, Cambridge
- Heiser J, Phan D, Agrawala M, Tversky B, Hanrahan P (2004) Identification and validation of cognitive design principles for automated generation of assembly instructions. *Proceedings of the working conference on advanced visual interfaces*. Gallipoli, Italy
- Hurley SM, Novick LR (2006) Context and structure: the nature of students' knowledge about three spatial diagram representations. *Think Reason* 12(3):281–308
- Larkin J, Simon H (1987) Why a diagram is (sometimes) worth 10,000 words. *Cogn Sci* 11:65–99
- Novick LR (2006) Understanding spatial diagram structure: an analysis of hierarchies, matrices, and networks. *Q J Exp Psychol* 59:1826–1856
- Novick LR, Hurley S, Francis M (1999) Evidence for abstract, schematic knowledge of three spatial diagram representations. *Mem Cognit* 27:288–308
- Taylor HA, Tversky B (1997) Indexing events in memory: evidence for index dominance. *Memory* 5:509–542

- Tversky B (1981) Distortions in memory for maps. *Cogn Psychol* 13:407–433
- Tversky B (2001) Spatial schemas in depictions. In: Gattis M (ed) *Spatial schemas and abstract thought*. MIT Press, Cambridge, MA, pp 79–112
- Tversky B, Lee PU (1998) How space structures language. In: Freksa C, Habel C, Wender KF (eds) *Spatial cognition: an interdisciplinary approach to representation and processing of spatial knowledge*. Springer, Berlin, pp 157–175
- Tversky B, Lee PU (1999) Pictorial and verbal tools for conveying routes. In: Freksa C, Mark DM (eds) *Spatial information theory: cognitive and computational foundations of geographic information science*. Springer, Berlin, pp 51–64
- Tversky B, Kugelmass S, Winter A (1991) Cross-cultural and developmental trends in graphic productions. *Cogn Psychol* 23:515–557
- Tversky B, Zacks J, Lee PU, Heiser J (2000) Lines, blobs, crosses, and arrows: diagrammatic communication with schematic figures. In: Anderson M, Cheng P, Haarslev V (eds) *Theory and application of diagrams*. Springer, Berlin, pp 221–230
- Tversky B, Agrawala M, Heiser J, Lee PU, Hanrahan P, Phan D, Stolte C, Daniel M-P (2006) Cognitive design principles for automated generation of visualizations. In: Allen GL (ed) *Applied spatial cognition: from research to cognitive technology*. Psychology Press, New York, pp 53–73
- Wagenaar WA (1986) My memory: a study of autobiographical memory over six years. *Cogn Psychol* 18:225–252
- Zacks J, Tversky B (1999) Bars and lines: a study of graphic communication. *Mem Cognit* 27(6):1073–1079
- Zwaan RA, Radvansky GA (1998) Situation models in language comprehension and memory. *Psychol Bull* 123:162–185