

Perception of Elementary Graphical Elements in Tabletop and Multi-Surface Environments

Daniel Wigdor^{1,2} Chia Shen¹ Clifton Forlines¹ Ravin Balakrishnan²

¹Mitsubishi Electric Research Labs
Cambridge, MA, USA
www.merl.com
forlines | shen @merl.com

²Department of Computer Science
University of Toronto
www.dgp.toronto.edu
dwigdor | ravin @dgp.toronto.edu

ABSTRACT

Information shown on a tabletop display can appear distorted when viewed by a seated user. Even worse, the impact of this distortion is different depending on the location of the information on the display. In this paper, we examine how this distortion affects the perception of the basic graphical elements of information visualization shown on displays at various angles. We first examine perception of these elements on a single display, and then compare this to perception across displays, in order to evaluate the effectiveness of various elements for use in a tabletop and multi-display environment. We found that the perception of some graphical elements is more robust to distortion than others. We then develop recommendations for building data visualizations for these environments.

Author Keywords

Visualization, tabletop, multi-display, multi-surface

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Table-centred computing environments have recently been the focus of much research [11, 21, 22, 28, 30, 31, 32, 33, 36, 37, 40, 41]. Recent efforts in combining tabletops with additional displays have resulted in display-rich environments, allowing for significantly larger information spaces. Research in this area has focused primarily on providing mechanisms for enabling multiple users to interact with information on these multiple displays [13, 20, 24, 37, 41]. Less attention has been paid to investigating how effective these different display angles and positions are for *presenting* information.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2007, April 28–May 3, 2007, San Jose, California, USA.
Copyright 2007 ACM 978-1-59593-593-9/07/0004...\$5.00.

Information visualization relies on encoding numerical values as graphical features. In order for data to be accurately decoded by readers, it is essential that they are able to precisely perceive the relative magnitudes of those features. Content displayed in tabletop groupware, however, is subject to at least two types of distortion. The first arises when multiple users sit around the table, and thus view content from different orientations [22, 23, 33, 40] (Figure 1, left). The second occurs when there is a large variance in the user’s viewing angle of different parts of the display (Figure 1, right), where objects displayed farther away from the user appear smaller and more susceptible to perspective effects than objects displayed closer to the user.

Humans clearly have some ability to compensate for these visual distortions: we are, for example, able to recognize objects under rotation [34], and perception of shapes drawn on a plane remain constant under rotation [27]. It is unclear, however, how much of an impact the second type of distortion – due to variance in viewing angle – might have on the accuracy of perception of the relative magnitudes of the elementary graphical elements [3, 5, 6] employed in information visualisations. In this paper, we examine this question in two experiments: one conducted on a single display oriented at varying angles on the continuum from vertical to tabletop, and the second on a dual tabletop *and* vertical display configuration. Our results indicate that some graphical elements are better suited than others for encoding information in tabletop and multi-surface environments.




Figure 1. Two types of distortion inherent in multi-user tabletop groupware. (Left) Information oriented “correctly” for one user is rotated for any other. (Right) Information located farther away from the user is subject to perspective distortion due to the differences in distance and viewing angle.

RELATED WORK

Research from three fields is relevant to our work. First, we briefly survey work in HCI which describe compelling applications for tabletop and multi-surface environments. Next, we examine work of theorists in the basic elements of information visualization. Finally, we survey research on perception of objects under 3D rotation, and the design of experiments involving magnitude perception.

Tabletop and Multi-Surface Environments

A typical tabletop system overlays a display and input device, creating a large, horizontal, direct-touch interface. Research in the area is generally divided into two camps: interaction design, [23, 31, 33], and psychological and sociological issues affecting the use of tabletops [14, 22, 30, 40]. There are also a number of projects which seek to augment tabletop interaction devices with vertical displays, [20, 21, 36, 37, 41]. Although advocated for use as information displays, we are unaware of any systematic examination of the effectiveness of tabletop and mixed-orientation display environments for the graphical display of quantitative information. In this work, we seek to inform the designers of visualizations for each of these distinct types of systems.

Basic Visual Elements

The basic elements of information visualization have been seminal and distinctly defined by two works: first, by Bertin [3], and again by Cleveland and McGill [5, 6]. Each of these authors believed that every visualization of information required that the data be encoded using one or more of the *visual variables* (Bertin) or *elementary perceptual tasks* (Cleveland & McGill) they described. Although each defined their own list of these, it is Cleveland and McGill's list (Figure 2) which is most relevant to our work. This is because they recognized that some of the broader categories of Bertin's list contained elements to which we have varying acuity, and so split those categories along perceptual lines. The effectiveness of their categorization was demonstrated by their experiments of perceptual ability, through which they were able to order their list in terms of how accurately a viewer can decode quantities represented using each visual variable [5].

An example of the splitting of Bertin's variables along perceptual lines is the dimensional breakdown of Bertin's *size* variable. Cleveland and McGill split Bertin's *size* between three elementary perceptual tasks: *length*, *area*, and *volume*. This was done because, as was found in a study by Baird [1], each is perceived with varying accuracy. Since we are concerned with users' perceptual ability in evaluating visual variables, we will use Cleveland and McGill's definitions, although our findings should be equally applicable to those seeking to apply Bertin's work to tabletop and multi-surface environments, and more generally to any designer of information visualizations.




Figure 2. Cleveland and McGill's elementary perceptual tasks. All visual representations of quantitative information require decoding using one or more of these. Image from [5].

It should be noted that each of Bertin's, and Cleveland and McGill's works offer a great deal more than the definition of these variables, including guidance for visualization designers wishing to leverage principles of perception, and describing other stages of perception of information visualization. Readers might wish to examine follow-up works by Green [19] and Carpendale [4], and applications of their results [2, 26].

3D and Angled Perception

The work of psychophysicists in the area of perception of objects at an angle and of 3D objects is vast, and providing a complete review of work on human perceptual mechanisms is beyond the scope of this paper. Some of the work from this area, however, is of particular interest.

It is likely that in order to perform magnitude perception of visual elements presented on an angled surface, some amount of mental rotation must be performed. Research on rotation of objects around the axis perpendicular to the display surface has shown that there is a linear relationship between the angle of rotation and the time required to perform it [7, 8, 34]. Also of interest is that all of gender, level of spatial ability, and visual field of presentation seemed to affect the speed of rotation, although they do not seem to affect accuracy [39]. Our desire to measure accuracy, rather than time, meant that we did not perform any screening for these in recruiting our participants.

Finally, psychophysicists have examined the issue of the effect of viewing angle on the perception of elementary graphical objects. In particular, the work of Schneider et al. [29] examined the effect of viewing angle on the perception of length. The angles of interest in their work, however, were the angle between the line and the center of the retina. We wished to examine in-situ perceptual accuracy, and so our experimental participants were free to move their eyes. As such, Schneider et al.'s results are not directly relevant.

Magnitude Perception

Models for the perception of relative magnitudes of a variety of physical phenomena provide an understanding of the internal perceptual processes in the brain [12, 35]. While valuable, the details of these models are beyond the scope of the present paper, and we refer the reader to Gescheider [16] for a thorough overview. Of more direct relevance to us is the ongoing debate as to how best to measure a participant's perceived magnitude of a given stimulus. Much of this work is similar in technique to the experiments performed by Cleveland and McGill [6]: a *modulus* object is shown to a participant who is then asked to evaluate the relative magnitude of subsequent objects as a fraction or ratio to that modulus. This technique, however, measures the *reported* relative magnitude, meaning that the accuracy of the collected data depends on participants correctly converting their perception into a numerical quantity. There is ongoing debate as to how best to collect the *actual* perceived magnitude, including applying models to results so as to avoid relying on participant reports [10, 25, 42]. In the present work, we avoid this controversy, since our goal is not to directly measure perceived magnitude, but rather the *change* in perception between multiple viewing angles. We require only that this change be present in whatever result we collect, and not that the reported results precisely match the actual perception. With this in mind, in designing our experiments, we have not attempted to model the underlying perception, as was done by Narrens [25], but have instead relied on the reported results, as was done by Cleveland and McGill [6].

We model our experimental design on that employed by Cleveland and McGill, with one important exception. In their work, they report that as the distance between modulus and stimulus objects increases, the error of the participants' estimate of magnitude increases [6]. What may be a flaw in their experimental design, however, is that distance is correlated with the order of magnitude comparison against a given modulus. More recent research has shown that subsequent comparisons against a single modulus are not generally independent, and so error tends to increase over time [16]. As a result, it may be the case that the effect for distance reported by Cleveland and McGill may actually be an effect of the ordering of comparisons. In our experiment, we have taken steps, described in the Task and Procedure section, to eliminate this confound.

Also relevant is the issue of the independence of magnitude estimations. Repeated studies have shown that changes in one, sufficiently different property do not affect the perception of a second visual property of that object [18]. This is especially relevant to the present study, where we will be comparing magnitude perceptions across various viewing angles of a display. This past work suggests that the brightness and colour representations of on-screen objects that might change with the viewing angle will not affect the perception of the magnitude of those objects.

TERMINOLOGY

Before discussing the details of our work, it is helpful to establish a set of terms for the various presentations of objects and imagery that we will use in our experiments. When discussing the orientation of the physical display, we will use the terms *vertical* for an upright display, and *tabletop* for a display laid flat. The display's orientation angle will range from 0° (tabletop) to 90° (vertical). The position of on-screen imagery will be defined by where it appears on the *left/right* (on-screen x-axis) and *up/down* (on-screen y-axis) axes. The orientation of on-screen imagery will be defined as *upright* when aligned with the up/down axis, and *lateral* when aligned with the left/right axis. Up/down and left/right distances between objects on an angled display are distinct: increases in up/down distance could increase the relative visual distortion between two objects depending on orientation of the display, while changes in left/right distance may not.

GOALS OF THE PRESENT WORK

Our ultimate goal is to provide guidelines for overcoming the visual distortion(s) introduced by working in a tabletop or multi-surface environment. We have thus extended the work of Cleveland and McGill, by conducting two experiments evaluating the performance of their elementary perceptual tasks in these environments. In the first, we sought to determine which elementary perceptual tasks are appropriate for use in a single display environment. We believed that, when the display is tilted close to or on the horizontal, the distortion introduced by the perspective change between two graphical elements displayed at different up/down distances from the viewer may impair users' abilities to perform an accurate estimate of their relative magnitudes. The second experiment extends the goals of the first by examining which basic graphical elements work best for comparisons across two displays: one oriented vertically, and the other oriented horizontally as a tabletop. We believed that since the difference in angle of the visual variables to the user is even greater between vertical and tabletop displays, so too would the level of impairment of perception be greater.

Tufte claims that visualizations of quantitative information must afford two levels of use [38]. First, a quick glance should provide an immediate, if coarse, understanding of a data set. Second, finer detail should be available upon closer scrutiny. It should be noted that, even when encoding information in such a way that it requires the execution of what we have found to be a less preferred elementary task, it is possible to provide additional visual information or cues that would allow the data to be decoded, such as by providing numeric values written beside each visual encoding. The goal of the present work, therefore, is not to claim that it is impossible for a reader to perceive values encoded with these variables, but rather to provide guidance for designers in selecting the variables which best facilitate the first of Tufte's requirements: a quick glance will provide coarse understanding and comparison of relative values.

EXPERIMENT 1: SINGLE DISPLAY

Goals and Hypotheses

In this experiment, we wished to examine whether compensatory processes within the brain would allow for uniform perception of visual variables across a tabletop surface. In order to evaluate this, we asked several participants to make magnitude comparisons of pairs of values of various visual variables on a display oriented at four different angles, and measured the accuracy of their estimates. Our hypotheses are as follows:

- H1: As the display is tilted, the accuracy of relative magnitude judgements decreases.
- H2: The up/down distance between objects is positively correlated with the increase in error in magnitude judgements due to display angle.
- H3: Different visual variable types have differing increases in the error in judgements.
- H4: Lateral presentations of objects experience less error in magnitude judgements due to display angle than upright presentations.
- H5: When the up/down on-screen positions of the modulus and stimulus objects are the same, judgement accuracy will be consistent across display angles.

The distinction between H2 and H5 is important: our belief is that, for comparisons where the amount of perspective distortion is consistent between the modulus and stimulus, there will be no increase in error as the display is angled, since the distortion will affect the two equally.

Additionally, as was previously stated, we hypothesize that the effect of distance on error reported by Cleveland and McGill may have been due to an experimental confound. As described in the Task and Procedure sections below, we have taken steps to eliminate this confound. Subsequently, we have an additional hypothesis:

- H6: There will be no effect for left/right distance on the accuracy of magnitude perception.

This contradicts the conclusions of Cleveland and McGill, but is based on recent results in magnitude perception [16].

Visual Variables

The visual variables in our experiment (Figure 3) are a subset of those examined by Cleveland and McGill [6], which differ from those in their earlier work [5] (Figure 2). Specifically, between these two publications, *curvature* was removed, and *direction* was replaced with *slope*. To maintain consistency, we based our list on that in their more recent publication [6]. We asked participants to report their perceived magnitude of each of *length*, *angle*, *position*, *slope*, and *area*. Each of *length*, *angle*, and *position* were presented in two orientations, one laterally and the other upright. This was done because we believed that lateral presentations would suffer less visual distortion when the display is laid flat, and so would present varying difficulty for decoding by the user.

Like Cleveland and McGill, we chose not to test the visual variable *colour saturation*. This choice was made for two reasons: first, because the presentation of colour to the retina is not distorted by varying viewing angle, and is therefore unlikely to be perceived differently on angled displays. Second, display technology is such that accurately reproducing colours across viewing angles is difficult, creating a likely confound. We also chose not to examine the *position, same scale* variable, since this would not lend itself to up/down position modulation, and because their results suggest that the differences between it and the *position, different scale* variable are not substantial.




Figure 3. The eight visual variables we presented in our experiments – a subset of those in Cleveland and McGill [6].

Apparatus

Participants were seated at a table, which was 66cm off the floor, on which was positioned a flat panel 43cm x 27cm NEC MultiSync LCD 2070WNX display with 1680x1050 pixel resolution. Each visual variable occupied a maximum on-screen size of 10cm². The display was in the “portrait” orientation relative to the user, so that it was taller than it was wide, and was attached to a mount which could reposition the display at 90° (vertical), 60°, 30°, or 0° (tabletop). To minimize the impact of display orientation, contrast and brightness were adjusted, and all imagery was displayed in black on a white background. When repositioned, the display was moved so that the top of the display, when angled as a tabletop, would be in the same location as the bottom of the display when vertical. This meant that the maximum distance that an object could appear on-screen from the user’s eye was approximately the same in these two conditions, so as to eliminate any confound between distance and viewing angle.

Participants were instructed to position their chair comfortably, and to keep it in the same position throughout the experiment. As was done in previous experiments [6, 40], the relative position and angle of the user’s head was not controlled, to allow them to perform whatever compensatory positioning of their head they felt was necessary, thus better approximating real-world conditions.

The apparatus was placed in a dark room, with the only light sources being the displays and a small light used to illuminate the keyboard, which was positioned carefully to ensure that it caused no glare on the display.

Task and Procedure

Our experiment was similar to Cleveland and McGill's, which presented a group of four numeric values encoded using a single visual variable type [6]. In their study, participants were asked to compare each of three of the objects with the fourth, largest, modulus object, by entering their magnitude as a percentage of the modulus'. As we have previously described multiple comparisons against a single modulus are not independent [10]. Researchers have suggested that each magnitude comparison can be made independent by forcing the participant to revisit the modulus, or by showing a new modulus for each stimulus [15, 17]. Although we mimicked all other aspects of Cleveland and McGill's experiment, we modified it such that a new, randomly determined modulus was presented for each randomly determined stimulus object for comparison. We used the relative placement of objects from their design [6], but repeated their pattern three times on the display, so that there were 3 modulus and 9 stimulus positions. This was done to increase up/down distances between moduli and stimuli for larger visual distortion. Figure 4 illustrates the locations of moduli and stimuli.




Figure 4. (Left). Experimental display in Cleveland and McGill [6]. (Right). Display used in the present experiment. We show possible modulus locations (M), and possible stimulus locations (S).

Participants were briefed on the task, and shown each of the 8 types of visual variables used in the study. To ensure understanding, they were asked to practice interactively with the experimenter: after reporting their response, they were told the correct response for each trial. During the actual experiment, no accuracy information was given.

Each participant was presented with a series of pairs of visual variables, and, were asked to report the magnitude of the stimulus as a percentage of the modulus [6]. The amount of time participants were allowed to view the stimulus was not limited, ensuring that user responses would be as accurate as possible [6]. When the participant began to type, the on-screen objects were hidden, replaced with their typed text. They could recall the objects for comparison by pressing *escape*; these trials were marked as erroneous. Each session lasted about 2 hours.

Participants

Twelve participants (7 men and 5 women) between the ages of 22 and 47 were recruited from the local community, and paid \$30 for their participation. Participants had completed at least one year of an undergraduate degree.

Design

Eight visual variables were included in the experiment, although each participant was presented with only 4 of them. The pairings were balanced so that, if a participant was to judge a particular variable, they would judge both the upright and lateral orientations of that variable. The remaining pairings of the visual variables was fully counter-balanced between participants

Three on-screen modulus positions were employed, along with 9 on-screen stimulus positions. For each visual variable and for each display angle, the participant completed three magnitude estimation trials: a total of 81 comparisons. These 81 comparisons were done in a block, so that the participant was not asked to jump back and forth between variable types. At the beginning of each block, on-screen instructions informed the participant of which visual variable would be presented, and they were required to complete 3 practice trials before beginning the block.

For each of the 4 display angles (90° vertical, 60°, 30°, and 0° tabletop), the participant completed one of these blocks of 81 trials for each of the four visual variables. The ordering of display angles, and the order of visual-variable types within each position, was controlled with a Latin-square. The experimental design can be summarized as:

12 participants x
 4 display angles x
 4 visual variables (per participant) x
 3 modulus positions x
 9 stimulus positions x
 3 magnitude estimates =
 15,552 total comparisons

Results

Our dependent variable was *error*, calculated as in [6]:

$$error = | \text{judged percent} - \text{true percent} |$$

In all cases, practice trials, and trials with a completion time or *error* of more than two standard deviations from the mean were excluded (131 trials). A computed independent variable was *up/down distance*, the distance between the modulus and stimulus objects along the vector perpendicular to the edge of the display closest to the user (always one of 0cm, 14cm, 28cm). This measure is independent of another computed variable: the *left/right distance* between those objects.

As was expressed in H1, we anticipated that as the display was tilted away from the user, and therefore as the visual distortion of the on-screen objects increased, we would find a corresponding increase in error. Looking only at the effect of display angle on error, we did see this effect: for each of 90°, 60°, 30°, and 0°, the mean for *error* across all visual variables was 8.9%, 8.3%, 10.5%, and 12.0%, $F_{3,208}=129.21$, $p < .0001$. The levels for each of 30° and 0° were significantly different from each other, and from 60° and 90°, which were not significantly different from one another. Thus, we confirm H1.

Although this overall difference is statistically significant, such a small increase in error between display angles for a first-stage analysis of a visualization is likely to be acceptable. However, as we break-down the source of this error, we see that the decrease in accuracy is actually much more significant in several cases. In particular, we are able to confirm H2 – that *up/down distance* is positively correlated with the significance of this increase in error. As we see in Figure 5, for the 30° and 0° display angles, the amount of error increased with *up/down distance*, while this error remained insignificantly changed for 60° and 90°. This suggests that the perspective distortion introduced by the display being angled has caused this difference.

In the 0°, or tabletop condition, the error increases from 8.69% for up/down aligned objects, to 12.61% for objects separated by approximately 14cm up/down, to 15.77% for objects separated by 28cm up/down.




Figure 5. Size of error by up/down distance between objects for each display angle.

We are also able to confirm H3: that some visual variables are more subject to this error than others. In all cases, *up/down distance* did not significantly affect error for the 90° and 60° angles. For the remaining angles, however, all visual variables had significant differences at the $p < 0.001$ level, in *error* between different levels of *up/down distance*, though to varying degrees. As we see in Figure 6, each of the visual variables was affected in how accurately it was perceived by the increase in perspective distortion. In Table 1, we see the precise mean error sizes for each.

We can partially confirm H4, that variables presented laterally would suffer less when presented on a tabletop display. For the three variables presented both upright and laterally, two showed a significant difference between upright and lateral presentation ($F_{3,207} = 102.7, p < 0.001$). As seen in Table 1, *length* was not significantly different, while perception of both *angle* and *position* was significantly better at lateral than at upright presentations.

We are also able to confirm H5, that in trials in which the modulus and stimulus are at the same up/down position, *error* will be consistent across display angles. For all visual variables, the level of *error* was not significantly different between any display angles for those trials where *up/down distance* between the modulus and stimulus objects was 0.

Finally, we are also able to confirm H6. For 7 of the 8 visual variables, we found that the left/right distance between the stimulus and modulus did not have a significant effect on error. The only exception to this was *slope*: for *slope* judgements, as *left/right distance* increased, so too did *error*. Post hoc testing revealed that each of these levels of *error* for *slope* was significantly different from one another (mean = 13.42, 14.08, 15.29) for each of the three left/right distances). That this effect is not present for the other variables directly contradicts the work of Cleveland and McGill [6].




Figure 6. Error size by up/down distance between compared objects for each visual variable type on a tabletop display.

Table 1. *error* (mean, standard deviation) by up/down distance for each visual variable type, displayed on a tabletop.

Up/down Dist.	Length (U)	Length (L)	Angle (U)	Angle (L)	Position (U)	Position (L)	Slope	Area
0cm	6.68, 6.28	7.34, 6.22	8.93, 7.76	8.68, 6.95	6.75, 6.70	7.71, 9.12	11.75, 9.94	11.64, 10.47
14cm	8.11, 5.38	7.72, 5.40	14.94, 10.07	12.24, 9.84	12.90, 6.58	11.76, 12.73	17.88, 10.69	15.57, 9.97
28cm	9.60, 6.52	9.24, 5.75	20.10, 9.96	14.02, 10.06	17.92, 9.29	12.53, 7.88	25.60, 12.12	17.31, 9.62

Discussion

Our results suggest that certain types of information visualizations may need to be modified in order to be displayed on a tabletop, since, as we have seen, comparisons of some basic elements of those visualizations may be distorted when working on the table. As such, a natural extension to the tabletop is to add vertical displays on which to show these visualizations, since comparisons of visual variables on the same, vertical surface, would not be subject to the distortions we have found for tabletops. Before advocating for a multi-surface environment, however, we wished to determine first how well users of such systems will be able to compare visual variables *between* tabletop and vertical surfaces.

EXPERIMENT 2: MULTI DISPLAY

Goals and Hypotheses

We extend our exploration from the single display study in Experiment 1 to an environment with both tabletop and vertical displays, such as those discussed in the literature [9, 13, 41]. We repeated a similar procedure from the first experiment, this time placing the modulus and stimulus objects on different displays kept immediately adjacent to one another. The display with the modulus object remained vertical, while the stimulus display was positioned either vertically (i.e., *vertical+vertical* condition) or as a tabletop (i.e., *vertical+tabletop* condition). Since this arrangement would result in an even greater difference in apparent angle between the two objects, we believed we would find a corresponding increase in the amount of error in the reported relative magnitudes. Our hypotheses were:

- H1: There is an increase in error when comparing visual variable magnitudes across vertical and tabletop displays (i.e., *vertical+tabletop*) versus comparing on displays of a single orientation (i.e., *vertical+vertical*).
- H2: The error increase when comparing between displays is unevenly distributed across visual variable types.
- H3: The size of the error on the *vertical+tabletop* condition is larger than in the *up/down distance = 28cm* condition on the tabletop in the previous experiment, since the angular distortion is greater.

Apparatus

The software and visual variable rendering were nearly identical to the first experiment. We changed the physical apparatus so that the modulus was shown on one display, while the stimulus variable was rendered on an adjacent display. Both were rendered on identical Dell 2000FP displays, with the same settings for brightness and contrast. As in the first experiment, the apparatus was installed in a darkened room with a small light for keyboard illumination.

Unlike in the first experiment, the left/right and up/down on-screen positions of the displayed objects remained constant. The up/down positions of the two variables were aligned, and the left/right position was such that they were presented adjacent to one another. The display with the stimulus was placed either vertical or as a tabletop. The display was positioned such that the axis of rotation was the up/down centre of the visual variables, effectively fixing the physical position of the centre the variables. Figure 7 illustrates display angle and the on-screen positions.




Figure 7. Display conditions for our second experiment. (Left) Vertical+vertical. (Right) Vertical+tabletop. The 3D position of the stimuli is fixed between conditions.

Task and Procedure

As in the first experiment, participants were briefed on the experimental task, and shown each of the 8 visual variables used in the study. To ensure proper understanding, they were asked to practice interactively with the experimenter, and, after reporting their response, were told the correct response for each trial. During the actual experiment, no accuracy information was provided to the participant.

Each participant was presented with a series of pairs of visual variables, and asked repeatedly to report the magnitude of the stimulus as a percentage of the magnitude of the modulus. When the participant began to type, the on-screen objects would be hidden, and replaced with their typed text. We recorded the value of their response.

Participants

8 male undergraduate students, aged 19 to 26, from the local university community participated in the experiment. All had normal or corrected-to-normal vision, and were given \$20 upon completion of their participation.

Design

Each participant was presented with both display conditions, with the ordering of the two balanced between participants. For each display condition, they were asked to make comparisons of all 8 types of visual variables, in blocks of 31 comparisons. The ordering of visual variable presentation was balanced with a Latin-square between display condition for the same participant and between participants. The study design can be summarized as:

8 participants x
2 display conditions x
8 visual variables x
31 magnitude estimates =
3,968 total comparisons

Results

As in the previous experiment, our dependent variable was *error*, calculated, as in Cleveland and McGill [6], as:

$$error = | \text{judged percent} - \text{true percent} |$$

In all cases, practice trials, trials in which the participant called for the objects after they began to type, and trials with a *time* or *error* of more than two standard deviations from the mean were excluded.

As in the first experiment, we are able to confirm H1, that there is an increase in error between values of a visual variable at the same display angle vs. comparing on displays of differing angles. In all cases, the difference between the *vertical+vertical* and the *vertical+tabletop* presentations is statistically significant at the $p < 0.001$ level. The results for each visual variable can be seen in Table 2.

As expected from the first experiment, we can also confirm H2, that the increase in *error* when comparing across displays (both in the *vertical+vertical* and *vertical+tabletop*) is unevenly distributed across visual variables.

Interestingly, we are unable to confirm H3, that the increase in *error* for the vertical+tabletop condition would be larger than in the first experiment's tabletop condition. Although the results for the vertical+vertical condition are nearly identical to that of the first experiment, the *error* size is, in many cases, actually lower for the vertical+tabletop condition in the second experiment than it is in the first experiment's tabletop condition. In only one case, when the *position* variable is presented upright, is the error higher in the second study than it is in the first.

Table 2. *error* (mean, and standard deviation) for both display conditions for each visual variable.

Display Condition	Length (U)	Length (L)	Angle (U)	Angle (L)	Position (U)	Position (L)	Slope	Area
Vertical+ Vertical	6.15, 5.06	6.02, 5.98	8.89, 8.01	7.73, 6.61	5.43, 6.55	5.45, 4.58	8.73, 7.67	10.32, 7.84
Vertical+ Tabletop	7.89, 5.99	7.13, 5.37	18.05, 8.65	15.06, 7.92	21.35, 11.03	12.26, 6.20	24.68, 12.10	17.43, 8.68

DESIGN RECOMMENDATIONS

From the results of these two studies, we are able to make several recommendations for the design of information visualizations in tabletop and multi-display environments.

Tabletop Visualization

For the design of graphical information visualizations for tabletop displays, we draw the following key results from our first experiment:

1. Perception of relative values of visual elements is less accurate on a tabletop when those presentations are not at the same up/down distance from the user.
2. The larger the up/down separation between elements on a tabletop, the less accurate comparisons between those elements will be.
3. Some visual variables are more accurately compared than others. The list, from most to least accurate, is: *length* (l), *length* (u), *position* (l), *angle* (l), *area*, *position* (u), *angle* (u), *slope*.
4. Of the visual variables which are more robust for tabletop display, two (*position*, *angle*) require that they be presented laterally relative to the viewer in order to maintain their robustness to reorientation.
5. If variables must be perceived both upright and laterally, *position* is more accurate than *angle*.
6. Larger left/right distances between compared elements to be do not yield an increase in errors, with the exception of the *slope* visual variable.

From these results, we can make several design recommendations. The first, second and sixth results above suggest that items to be compared should be carefully

positioned at the same up/down distance from the user, in order to equalize the amount of perspective distortion. With multiple users, it may be necessary to repeat information displays, or position them within the private area [30] of the particular user for whom they are intended. These approaches would also conform with the fourth result, by avoiding undesirable orientation of the visual variables that are rotationally sensitive.

From the third result, we see that information visualizations that encode using length, distance, and angle are significantly less subject to error than those that make use of slope or area. From the fourth result, however, we see that, for information visualizations that encode two dimensions of information, position is not an appropriate variable, since upright and lateral positions would effectively be subjected to different scales. This is seen in the graphs in Figure 8.




Figure 8. Values encoded in position (left) are less accurately perceived on a tabletop than those encoded in length (right).

For visualizations which attempt to use multiple visual variables to simultaneously encode multiple categories of data, it is especially important to present objects for comparison at the same up/down distance. An example of such a graph is a box plot, which use both *position* and *length* to encode information. When comparing multiple box plots on a tabletop, those displayed farther away from the user can be compared for their lengths, but not for the up/down position, to those positioned closer on the table. As such, the size of inter-quartile ranges can be compared across multiple plots, but the actual bounds of that range will be inaccurately compared. Figure 9 illustrates an alternative design, which replaces the use of *position* with *length* judgments, creating a visualization more suitable for a tabletop display.




Figure 9. This box plot design replaces the use of *position* with *length* for illustrating the bounds of the inter-quartile range.

Finally, our findings for the accuracy of comparisons of *slope* and *area* values suggest strongly that they should not be employed on the tabletop. This is consistent with previous recommendations made by Cleveland and McGill, who suggest that every line graph should be accompanied with a second graph encoding the slope values of the first [6]. Our findings match theirs, except that this second graph would best be designed to use length rather than position judgments, such as a bar graph.

The examples shown here are just a few of the implications of the application of our results to visualization design for the tabletop. Those seeking to design tabletop systems might be guided by these examples, as well as by seeking to apply our results to their own visualization designs.

Mixed-Display Visualization

In our second experiment, we found that errors of comparison of visual variables across displays were not as severe as those made directly across large up/down distances on a tabletop. We also found that one of the visual variables, *position*, is even less robust to differing orientations between displays than it is to large perspective distortions on the same display.

Designers or systems which mix tabletop and vertical displays will need to apply the tabletop design results outlined above when designing for their tabletop surface alone. When designing software that is meant to be used on multiple surfaces simultaneously, the following findings of our second study may be useful:

1. Information visualizations should not be compared across display orientations.
2. If comparisons across display orientations are unavoidable, the first three in this ordered list by accuracy are far better than the rest: *length (l)*, *length (u)*, *position (l)*, *angle (l)*, *area*, *angle (u)*, *position (u)*, *slope*.
3. If a second variable is required to be perceived both upright and laterally, *angle* is more accurate than *position*, unlike on a tabletop.

The key difference between designing information visualizations for tabletop and multi-surface environments is that, in the latter, *position* is even less accurately perceived than *angle* for upright presentations. Given this, in multi-surface environments where the variables will be seen both upright and laterally, *angle* should be used before *position* for encoding information.

DISCUSSION AND FUTURE WORK

One of the more interesting results from our first experiment was a contradiction with Cleveland and McGill [6], in that we found that left/right distance between compared objects did not increase error. As we have discussed, we attribute this finding to our application of findings in magnitude perception to our experimental design, and suspect that the previous finding was due to experimental confound.

What is probably the most surprising result of the second experiment is that, for the tabletop condition, there is no increase in error over the results of the first experiment, despite similar results for the vertical condition. As with any result in these types of experiments, accurately attributing this finding to an underlying cause is challenging. There are several possible explanations, and is a likely candidate for future study. Also a candidate for future study is extending this work to multiple vertical display orientations, very large displays, or volumetric displays, all of which may cause differing perspective distortion under certain circumstances. An important addition to this work will be the effects of rotation of objects around the axis perpendicular to the display surface, since this effect is also present in collaborative tabletop groupware (Figure 1). Ultimately, the development of a model for predicting the effects of all of these variables on perception accuracy would be of the most use to designers.

In both experiments, because we did not wish to create artificial limits on performance, participants were allowed as much time as they wished to perform the magnitude estimates. Examining a potential viewing time vs. accuracy trade-off is an avenue for further exploration.

Finally, many of the results reported in the two experiments are arguably small differences in perception accuracy. Interesting to examine in the future is whether or not these small differences in perception of visual variables can be overcome by providing additional contextual information in information visualizations that employ them. For example, by displaying a background grid behind a histogram, judgements of *location* may become de facto judgements of *length*, and therefore more robust to display on a tabletop or across display orientations. This, and other techniques should be considered by designers. The intention of this paper is to highlight an issue that needs to be overcome, not to attempt to close doors for future visualization designs.

ACKNOWLEDGEMENTS

We thank John Barnwell and John Hancock for technical assistance, members of the MERL and DGP labs for valuable discussion, and our experimental participants.

REFERENCES

1. Baird, J.C. (1970), *Psychophysical Analysis of Visual Space*, New York: Pergamon Press.
2. Beattie, V., Jones, M.J. (2002). The impact of graph slope on rate of change judgements in corporate reports. *ABACUS* 38 (2). p. 177-199.
3. Bertin, J. (1977). *La Graphique et le traitement graphique de l'information*. Flammarion, Paris, France.
4. Carpendale, M.S.T. (2003). Considering Visual Variables as a Basis for Information Visualisation. *Computer Science TR #2001-693-16*. Dept. Computer Science, University of Calgary.
5. Cleveland, W.S., McGill, R. (1984). *Graphical Perception: Theory, Experimentation, and Application to the Development*

- of Graphical Methods. *Journal of the American Statistical Association*, 79(387). p. 531-553.
6. Cleveland, W.S., McGill, R. (1985). Graphical perception and graphical methods for analyzing and presenting scientific data. *Science*, 229(4716). p. 828-833.
 7. Corballis, M.C., McLaren, R. (1984). Minding one's Ps and Qs: Mental rotation and mirror-image discrimination. *Journal of Experimental Psychology: Human Perception and Performance*, 10. p. 318-327.
 8. Corballis, M.C., Macadie, L., Beale, I.L. (1985). Mental rotation and visual laterality in normal and reading disabled children. *Cortex*, 21. p. 225-236.
 9. Coutaz, J., Lachenal, C., Dupuy-Chessa, S. (2003). Ontology for Multi-surface Interaction. *Interact'03*. p. 447-454.
 10. Ellermeier, W., Faulhammer, G. (2000). Empirical evaluation of axioms fundamental to Stevens's ratio-scaling approach: I. Loudness production. *Perception & Psychophysics*, 62. p. 1505-1511.
 11. Everitt, K., Shen, C., Ryall, K., Forlines, C., MultiSpace: Enabling Electronic Document Micro-mobility in Table-Centric, Multi-Device Environments. *TableTop 2006*. p. 27-34.
 12. Fechner, G.T. (1860). *Element der Psychophysik*. Leipzig: Breitkopf & Harterl.
 13. Forlines, C., Esenther, A., Shen, C., Wigdor, D., and Ryall, K. (2006). Multi-user, multi-display interaction with a single-user, single-display geospatial application. *UIST 2006*. p. 273-276.
 14. Forlines, C., Shen, C., Wigdor, D., Balakrishnan, R. (2006). Exploring the effects of group size and display configuration on visual search. *CSCW '06*. p. 11-20.
 15. Galanter, E. (2004), Extracting magnitude estimations of loudness from pairwise judgments. *The Journal of the Acoustical Society of America*, 115(5). p. 2534.
 16. Gescheider, G.A. (1988). Psychophysical Scaling. *Annual Review of Psychology*, 39. p. 159-200.
 17. Green, D.M., Luce, R.D. (1974). Variability of magnitude estimates: a timing theory analysis. *Perception & Psychophysics*, 15(2). p. 291-300.
 18. Green, M. (1991) Visual search, visual streams, and visual architectures. *Perception and Psychophysics*, 50. p. 388-403.
 19. Green, M. (1996). *Toward a Perceptual Science of Multidimensional Data Visualization: Bertin and Beyond*. <http://web.archive.org/web/19991104170601/http://www.ergogero.com/dataviz/dviz0.html>. (archive).
 20. Johanson, B., Hutchins, G., Winograd, T., and Stone, M. (2002). PointRight: experience with flexible input redirection in interactive workspaces. *UIST 2002*. p. 227-234.
 21. Koike, H., Nagashima, S., Nakanishi, Y., & Yoichi Sato. (2004). EnhancedTable: Supporting small meetings in ubiquitous and augmented environment. *PCM2004*. p. 97-104.
 22. Kruger, R., Carpendale, M.S.T., Scott, S.D., Greenberg, S. (2004). Roles of Orientation in Tabletop Collaboration: Comprehension, Coordination and Communication. *Journal of CSCW*, 13(5-6), p. 501-537.
 23. Morris, M.R., Paepcke, A., Winograd, T., and Stamberger, J. TeamTag: Exploring Centralized versus Replicated Controls for Co-located Tabletop Groupware. *CHI 2006*. p. 1273-1282.
 24. Nacenta, M. A., Sallam, S., Champoux, B., Subramanian, S., Gutwin, C. (2006). Perspective cursor: perspective-based interaction for multi-display environments. *CHI2006*. p. 289-298.
 25. Narens, L. (1996). A theory of magnitude estimation. *Journal of Mathematical Psychology*, 40. p. 109-129.
 26. Nowell, L.T. (1997). *Graphical encoding for information visualization: using icon color, shape, and size to convey nominal and quantitative data*. PhD thesis, Department of Computer Science, Virginia Polytechnic Institute and State University.
 27. Pizlo Z. (1994). A theory of shape constancy based on perspective invariants. *Vision Research* 34. p. 1637-1658.
 28. Rekimoto, J. (2002). SmartSkin: an infrastructure for freehand manipulation on interactive surfaces. *CHI2002*. p. 113-120.
 29. Schneider, B., Erlich, D.J., Stein, R., Flaum, M., Mangel, M. (1978). Changes in the apparent lengths of lines as a function of degree of retinal eccentricity. *Perception*, 7. p. 215-223.
 30. Scott, Stacey D. (2005). *Territoriality in Collaborative Tabletop Workspaces*. Ph.D. thesis, Department of Computer Science, University of Calgary.
 31. Scott, S., Carpendale, S., & Habelski, S. (2005). Storage bins: Mobile storage for collaborative tabletop displays. *IEEE CG&A*, July/August. p. 58-65.
 32. Shen, C., Ryall, K., Forlines, C., Esenther, A., Vernier, F., Everitt, K., Wu, M., Wigdor, D., Morris, M.R., Hancock, M., Tse, E. (2006). Informing the design of direct-touch tabletops. *IEEE CG&A*, 26(5). p. 36-46.
 33. Shen, C., Vernier, F., Forlines, C., & Ringel, M. (2004). DiamondSpin: An extensible toolkit for around the table interaction. *CHI2004*. p.167-174.
 34. Shepard, R.N., Metzler, J. (1971). Mental rotation of three dimensional objects. *Science*, 171. p. 701-703.
 35. Stevens, S.S. (1957). On the Psychophysical law. *Psychology Review*, 64. p. 153-181.
 36. Streitz, N., Geißler, J., Holmer, T., Konomi, S.i., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., & Steinmetz, R. (1999). i-LAND: an interactive landscape for creativity and innovation. *CHI.1999*. p. 120-127.
 37. Streitz, N., Prante, C., Muller-Tomfelde, C., Tandler, P., and Magerkurth, C. (2002). Roomware©: the second generation. *Ext. Abs. CHI 2002*. p. 506-507.
 38. Tufte, E. (1983). *Visual Display of Quantitative Information*. Graphics Press, Cheshire, Connecticut.
 39. Voyer, D., Brydon, M.P. (1990). Gender, level of spatial ability, and lateralization of mental rotation. *Brain and Cognition*, 13(1). p.18-29.
 40. Wigdor, D., Balakrishnan, R. (2005). Empirical Investigation into the Effect of Orientation on Text Readability in Tabletop Displays. *ECSCW2005*. p. 205-224.
 41. Wigdor, D., Shen, C., Forlines, C., Balakrishnan, R. (2006). Table-Centric Interactive Spaces for Real-Time Collaboration. *AVI 2006*. p. 103-107.
 42. Zimmer, K. (2005). Examining the validity of numerical ratios in loudness fractionation. *Perception & Psychophysics*, 67. p. 569-57.