

Empirical Investigation into the Effect of Orientation on Text Readability in Tabletop Displays

Daniel Wigdor, Ravin Balakrishnan

Department of Computer Science, University of Toronto

dwigdor | ravin @dgp.toronto.edu

Tabletop collaborative groupware is a newly re-emerging field in CSCW. The use of a tabletop display presents a unique challenge to interface designers: how to optimally orient displayed objects for viewing and manipulation by users situated at various locations around the table. A great deal of CSCW research has been conducted under the implicit assumption that textual elements should be oriented directly toward the reader, despite research that demonstrates that a simple, straight-on orientation is not necessarily ideal in all circumstances. Absent from this ongoing research dialogue, however, has been an empirical examination of user performance of reading text on tabletop displays at non-zero orientations. In this paper, we present two studies which examine the effect of text orientation on common tasks: the reading of a small piece of text, and the serial search for a label. We found that, though statistically significant, the effects of orientation on the performance of these tasks were less dramatic than might have previously been assumed. From this, we hope to help guide collaborative groupware designers as to when orientation should be “corrected”.

Introduction

In recent years, a great deal of research has been conducted into the design and implementation of tabletop interactive systems; for example, the works of Ringel et al. (2004); Rogers et al. (2004); Shen et al. (2003); Shen et al. (2004); Streitz et al. (1999); Streitz, et al. (2002); and Wu and Balakrishnan (2003). Especially exciting about this domain is that a tabletop application can be targeted to simultaneous use by multiple users seated around the table. This organization, however, creates a problem unique to the domain: how should on-display objects be oriented?

Research has been conducted into how participants in non-computer based, table-centered collaborative tasks make use of orientation (Tang 1991; Kruger et al. 2003). These investigations have shown that users prefer a straight-on orientation for reading text, and orient objects towards themselves, others, or in-line with shared artifacts to ease reading in different circumstances. They have discovered, however, that in a collaborative setting, a straight-on orientation toward the reader is not *always* desired or exercised. In fact, orientation is employed as a tool to aid in interaction with other users. Despite this, it seems that designers of collaborative research systems have, in general, opted to attempt to orient text towards the reader, as seen in Bruijn et al. (2001), Rekimoto et al. (1999), Shen et al. (2003), Shen et al. (2004), and Streitz et al. (1999). Thus, there is a tension between a desire to allow for the use of orientation as an aide to collaboration, and the designers' assumption that users need to have textual elements oriented towards them.

Absent from this previous work is a thorough investigation into the parameters for determining *when* a solution to the text orientation problem should be applied in the context of tabletop groupware. Although users seem to prefer a "normal" orientation of text (Kruger et al. 2003), and studies in the psychology literature in non-tabletop situations with constrained user head and body movement by Tinker (1972) and Koriat and Norman (1984, 1985) indicate that readability is compromised when text is oriented, is it possible that there are circumstances where it might be appropriate to ignore this preference in favour of a less preferred orientation that may aid collaboration in other ways? Is orientation so critical for text readability that design elements must be sacrificed in order for it to be addressed? Without empirical data quantifying the extents to which readability is compromised in less-preferred orientations on tabletops, it is difficult to make informed choices when confronted with these tradeoffs. Our present work provides such empirical data via two experiments that examine the effect of text orientation on the performance of tasks common to tabletop collaborative groupware. Based on the results, we hope to provide insights for system designers as to when the issue of text orientation takes precedence, and when it can be safely ignored. Although very infrequently done in the field of human-computer interaction, replication and extension of experimental work is an important aspect of research. Our work also contributes in this regard by re-examining and extending the studies of readability of text orientation by Tinker (1972) and Koriat and Norman (1985) to tasks relevant to the new domain of tabletop displays, and where users' head and body movements are unconstrained.

Throughout this paper we will be referring to various textual rotations. To help orient the reader, we will refer always to anticlockwise rotations of text. Figure 1 demonstrates the various orientations used in our experiments.




Figure 1. Degrees of rotation as referred to throughout this paper. In all cases, angles are measured relative to the edge of the table at which the participant was seated.

Related Work

Of relevance to the present work are two areas of research: first, an examination of how the issue of reading orientation has been addressed by researchers of tabletop collaborative groupware will provide context for the present research. Second, research in the field of cognitive psychology concerned with reading, especially of text at non-horizontal orientations, is reviewed.

Orientation in Collaborative Groupware Research

An examination of the role of artifact orientation in collaborative settings has been presented by Tang (1991) and Kruger et al. (2003). The lessons they present are clear: orientation has a role to play in collaboration, and any simple, one-sized-fits-all algorithmic solution to the issue of orientation of screen artifacts will deprive participants of some of the richness of interaction afforded in paper-based tasks. What they also provide is a clear demonstration of the intuitive notion that users regard text presented in a “normal” orientation as more readable. They both agree that, similarly to what was reported by Fitzmaurice et al. (1999), users do not always orient artifacts in the same right-angled orientation found in most modern GUI. For example, orientation for readability is often used as a mark of ownership, and re-orientation towards a collaborator is a strategy for passing-over or signifying a desire to share.

From this research, we make two important observations: 1. right-way orientation of artifacts is an important theme in tabletop collaboration, and 2. that users regard the orientation of task objects as important cues for ownership and collaborative behaviour.

Despite the tradeoffs inherent in the above observations, systems continue to be designed that attempt to “solve” the issue of text orientation. For example, by taking

advantage of advanced display technologies, in both Agrawala et al. (1997) and Matsushita et al. (2004), users view a common scene orientation, but textual object labels are oriented toward each participant. Additionally, many systems attempt to dynamically and automatically re-orient objects; Kruger et al. (2003) present a thorough review of these algorithmic approaches, to which we refer the reader for more information. The reasoning behind these approaches seems sound: if text is easier to read at a particular orientation, an ideal approach would always present that orientation to any user attempting to read the content. It is their implicit contention that right-way up reading is so important that the group-dynamic affordances described in previous research should be sacrificed. But, how real is the tension between readability and system flexibility?

In an attempt to answer this, we now examine the research in text orientation that has been conducted in the field of cognitive psychology.

Human Ability to Read Text at Various Orientations

Research into the effect of graphical considerations on reading performance began with Huey (1898). Although the issue of orientation was not discussed, the effect of vertical vs horizontal alignment of words was examined, as well as the effect of partially occluding portions of characters. The issue of orientation and its effect on reading was first explored in detail by Tinker (1972). He conducted two experiments: in the first, subjects performed the Chapman Cook Speed of Reading Test at various orientations (Chapman 1923). The Chapman-Cook test involves the presentation of a paragraph, where a single word “spoils the meaning” of the text. The subject is required to identify this word, either by crossing it out, or speaking it aloud. The speed of reading is measured by how quickly this is accomplished. In this experiment, Tinker (1972) found that the reading of text rotated at 45° in either direction was, on average, 52% slower than reading normally oriented text, and text rotated at 90° in either direction was 205% slower on average. In the second experiment, participants performed the Luckiesh-Moss Visibility Meter to determine the physical visibility of rotated text (Luckiesh 1944). Tinker (1972) discovered that the speed of reading was affected much more dramatically by orientation than was visibility, and thus concluded that visibility was not the only factor that contributed to the decreased speed of reading at non-standard orientations.

Koriat and Norman (1984) used a text-reading task to evaluate the relative merits of two theories of how mental rotation is performed. Although this was informative to the present study, it was their later work (1985) which examined the issue of readability of rotated text in more detail. Specifically, they examined the effect of rotating text on the identification of strings as words or non-words. They found that performance of their task in the range of -60° to 60° degrees from the horizontal was not significantly variable, but that a performance cliff was reached once rotation exceeded 60° in either direction. Once this cliff was reached, word/non-word identification speed decreased by more than 120% (Koriat and Norman 1985).

The work of both groups seems to confirm the intuitive notion that reading orientation should be a key concern to designers of tabletop systems. However: their work is not

directly applicable to tabletop collaborative groupware research. First, rather than allow participants free and natural movement during the experiments, the position and orientation of their heads was constrained. Second, in both sets of experiments, readability was determined by how quickly non-conforming strings were identified. In the Tinker (1972) study, this was done at the semantic level, as participants were required to find the word in paragraphs that spoiled their meaning. In the Koriat and Norman (1985) experiment, this was done at the syntactic level, as the study consisted of the presentation to the participant a series of strings of characters, which subjects were required to identify as either words or non-words. Though this experimental design enabled them to answer their research questions, we note that the identification of non-conforming or gibberish strings is not directly applicable to real user interface scenarios. In most applications, textual artifacts consist of either common words or domain terms that might be expected by the user. This assumption might aid in the reading of text at varying orientations, and so should be considered when evaluating user performance of reading at varying orientations.

In our studies, we attempted to provide a more “natural” environment where user head movements are not constrained, and measured the performance of reading non-gibberish text at various orientations. It was our hypothesis that, given this environment, the effect of orientation on task-performance would be less dramatic. If this were the case, we believe that the tension between orientation as a tool for collaboration and the apparent need to use the “right” orientation for text readability can be relaxed, and systems could begin to be designed that heed the observations of Tang (1991) and Kruger et al. (2003).

Experiment 1

Goals

Although we were certain that orientation would have a statistically significant effect on text reading performance (based on the results in the literature), we wished to measure the strength of the effect at each orientation. We also wanted to conduct an experiment that would measure performance in an actual collaborative tabletop environment as compared to the artificially constrained environments used by Tinker (1972) and Koriat and Norman (1985). To this end, we present the text to the user on a tabletop, and allow participants free movement of their body and head, to allow for any tendency toward personal re-orientation. We also wished to examine how the type of information presented might affect performance. We presented three types of stimuli: a single word, a short phrase, and a 6-digit number. In order to ensure that participants could easily comprehend the words and phrases, single word stimuli were only simple, 5-6 letter words, while the phrases presented were coherent and meaningful.

It was our hypothesis that the effect of orientation on reading of a single word would be less dramatic than what Koriat and Norman (1985) observed. Because our set consisted of only common words, participants would be able to trust their immediate

identification of a rotated word, rather than graphically examine each character, as was required in the Koriat and Norman (1985) experiment. Further, we believed that the performance in reading longer phrases would be better than what Tinker (1972) reported. Because our phrase set consisted only of short, logical phrases, participants would be able to rely on the context of the surrounding text to aid in identification of harder to read words. Lastly, we believed that the reading of numbers would be most affected by orientation, since no particular grouping of the numbers could be assumed. Thus, we expected that our results in this situation would be similar to that previously reported.

Apparatus

Text was presented to the user using a ceiling-mounted digital projector, and was projected onto a DIAMONDTOUCH tabletop, introduced by Dietz and Leigh (2001), at which the participant was seated. Although the DIAMONDTOUCH is intended for touch-input, we used it only as a display screen since object manipulation was not required in our experiment. While we could have just as easily projected the image onto any tabletop surface, we chose the DiamondTouch since it is one of the common platforms for tabletop research and has a diffuse surface that results in high-quality imagery. The text was presented in a sans-serif font, and rotated in software to be presented to the participant. There was no apparent degradation in the quality of the rendering at the various orientations.

Text entry was facilitated by a standard QWERTY keyboard placed directly in front of the participant, who was seated at a chair placed directly in front of and centered on the longer side of the table. The system was driven by a windows-based Intel Pentium 3.0GHz system, equipped with an NVIDIA GeForce FX Go5700. Figure 2 illustrates the experimental environment.




Figure 2. Top-down diagrammatic view of the experimental apparatus.

Participants

Fifteen participants, recruited from the university community, volunteered for the experiment. 12 were male, 3 were female. 13 spoke English natively, while the remaining 2 reported to be excellent readers. All self reported as excellent typists. Participants received no compensation for their participation.

Procedure

Users were repeatedly presented with a string which they were asked to read, memorize, and type-in to the system. We wished to measure how long they would spend reading the text before they were sufficiently confident that they would remember it long enough to type-in to the system.

The strings consisted of 5-6 letter words, 6-digit numbers, and 6-7 word phrases. The words and numbers were randomly selected and generated, while the phrases were selected from the corpus developed by MacKenzie and Soukoreff (2003). The location of each string on the screen was primed with the display of a red cross for 0.65 seconds before the appearance of the text. When the user began to type, the text would disappear from the screen, but would return whenever “escape” was pressed. Figure 3 illustrates the procedure:



Figure 3. (a) Orienting crosshair primes the participant as to the location of text. (b) After 0.65 seconds, the crosshair disappears and is replaced by the rotated text. (c) As the subject begins to type, the text disappears from the screen. The black arrow is for illustration only and was not present in the actual experiment.

Timing began when the string was first displayed, and stopped when the user began to enter input. In order to prevent participants from “racing through the experiment”, they were required to enter perfect input, and were told to correct their entry if it was incorrect when they pressed “enter”. Trials where such correction occurred were not included in our timing analysis. Participants were given initial instructions as follows:

The experiment will require you to read some text on screen. Each time text is presented, you should memorize it (as quickly as you can), and then start to type it into the computer. When you begin to type, the text will disappear from the screen.

At any time, you can view the text again by pressing “escape”.

You will be required to enter the text absolutely correctly. If you make a mistake, the system will tell you, and ask you to correct your entry. Press enter to begin the experiment.

Before each priming/string was presented, they were given the following on-screen instruction:

You will now be presented with a red cross, which will shortly be replaced with text. Read the text, keeping in mind that you will need to remember the text long enough to type it into the computer after it disappears. When you are ready, type the text in to the computer. If you make a mistake the system will alert you. Press "space" to begin.

To ensure that participants understood what was required, they were allowed a practice dataset consisting of several short phrases. They were instructed to enter as many as was required to become familiar with the apparatus. Participants were directed to rest as required between strings, but to continue as quickly as possible once a phrase was presented on-screen.

Design

Strings were presented in three datasets of 96 elements each: one set of 5-6 letter words, one set of 6 digit numbers, and one set of 6-7 word phrases. The order of presentation of the datasets was counterbalanced between-participants using a latin-square. Strings were presented on the screen centered at one of the four corners of the display, and in one of 8 orientations, including straight-up, and each position around the compass at 45 degree increments. Within each dataset the position and orientation of a string was randomized, but controlled such that each position and orientation at that position was presented an equal number of times. The content of each dataset was fixed, such that all participants entered the same strings in the same order. In summary, the design was as follows:

3 datasets (single word, number, short phrase) X
4 on-screen positions (each corner of the tabletop) X
8 orientations (starting at 0°, in 45° increments) X
3 strings at each position/orientation X
15 participants
= 4320 strings entered in total.

All orientations were measured relative to the side of the table at which the participant was seated. Because head and chair positions were not constrained beyond the need to reach the keyboard to enter text, the exact orientation of the presented strings relative to the participants' eyes was not measured. Assuming a comfortable typing distance, as was generally observed, the angle of the 0-degree oriented-text relative to the centre of the participant's nose was approximately 12° for the upper-quadrant cases, and 17° for the lower-quadrant cases, well within the range of orientations shown to have little effect on reading speed in the works discussed previously.

Alternative Experimental Designs

Two alternative experimental designs were considered, but rejected in favour of the design we have presented. We will briefly present each of the designs and explain why we rejected them:

Stroop

The Stroop test is one of the most famous in cognitive psychology. As first demonstrated by Stroop (1935), it was found that participants were slower to identify the color of a string if the text was an interfering color name than if it was not the name of a color. It is believed that the participants were reading the text more quickly than they were identifying the color, and that the string was interfering with the identification of the color.

Our first experimental design consisted of a modified Stroop test, where the strings would be presented in various orientations. We believed that the Stroop effect would continue to assert itself, and so demonstrate that the participant was able to quickly read the text. We rejected this design for two reasons: first, we wished to measure performance of reading of various types of strings. Second, we realized that, even if the effect continued to assert itself, that we would be demonstrating only that reading speed was above the color-identification threshold. While informative, this would fail to measure with sufficient fidelity the effect of orientation on reading performance.

Read Aloud

A read aloud design would have consisted of the presentation of text at various orientations, the user reading the string aloud, and the measurement of the total time required to read the text. This design would have mimicked the design presented by Huey (1898). We chose not to employ this design for two reasons: first, limitations in speech recognition technology would limit our reporting accuracy. Second, differences between silent and spoken reading speed would limit the efficacy of the experiment in demonstrating reading performance at varying orientation.

Results

We discarded from the timing data all trials where the participant entered erroneous data which was then subsequently corrected. Orientation did not have a significant effect on error rate ($F_{7,26} = 1.40$, $p = .1990$). The type of stimuli did have an effect on errors ($F_{2,26} = 34.04$, $p < .0001$), with mean error rates of 15%, 8%, and 4% for single word, short phrase, and number treatment respectively.

As expected, orientation had a statistically significant effect on speed of reading: single word ($F_{7,10} = 28.0$, $p < .0001$), short phrase ($F_{7,10} = 64.28$, $p < .0001$), and numbers ($F_{7,10} = 7.76$, $p < .0001$). Pairwise means comparisons of entry time of all types of text indicated that orientations of -45° , 0° , and 45° were not significantly different from one another, but were significantly different than the rest. In the single word and short phrase

conditions, time to entry of stimuli presented at orientations of -135° and 135° were not significantly different from one another, but were significantly different from those at -90° and 90° . In the number condition, time to entry of stimuli presented at all orientations beyond -45° and 45° were not significantly different from one another. The location of presentation of the stimuli among the 4 on-screen position did not significantly affect performance time ($F_{3,26} = 1.95$, $p = .12$).

As we hypothesized, the effect of orientation on reading speed was far less dramatic than had been previously reported. Contrary to our hypothesis, however, the effect of rotation on 6-digit numbers was the least dramatic of the three conditions. Table I summarizes the mean rates for each orientation and condition.

	<i>Single Word</i>			<i>Short Phrase</i>			<i>6-Digit Number</i>		
	μ (secs)	σ	% off 0°	μ (secs)	σ	% off 0°	μ (secs)	σ	% off 0°
-135°	1.19	0.67	64.70	3.82	1.52	107.13	2.85	1.06	17.48
-90°	0.92	0.40	26.60	2.66	1.07	44.25	2.85	1.63	17.19
-45°	0.78	0.60	7.98	2.07	1.02	12.62	2.36	1.17	-2.71
0°	0.72	0.22	-	1.84	0.86	-	2.43	1.57	-
45°	0.77	0.22	5.93	1.97	0.70	7.19	2.39	1.19	-1.65
90°	0.91	0.37	25.78	3.09	1.30	67.71	2.78	1.21	14.56
135°	1.35	1.00	86.42	3.90	1.97	112.82	3.01	1.24	24.26
180°	1.11	0.57	53.67	3.69	1.89	100.27	3.03	1.16	24.87

Table I. Summary of mean (μ) and variance (σ), of reading times, and percentage deviation from mean reading time for un-rotated text of the same type for each orientation, and each of the three conditions. Excludes trials where erroneous data was entered and subsequently corrected.

Figures 4-6 are boxplots for time required to read the stimuli under each condition, broken down by screen-quadrant. Our results correspond with those of Koriat and Norman (1985) in that, for the word and phrase conditions, the worst mean performance for reading stimuli in the upper quadrants of the screen was 135° in the upper-left quadrant, and -135° in the upper-right quadrant. Since the participant was positioned at roughly the centre of the table, these orientations represent text written bottom-to-top, and away from the participant.

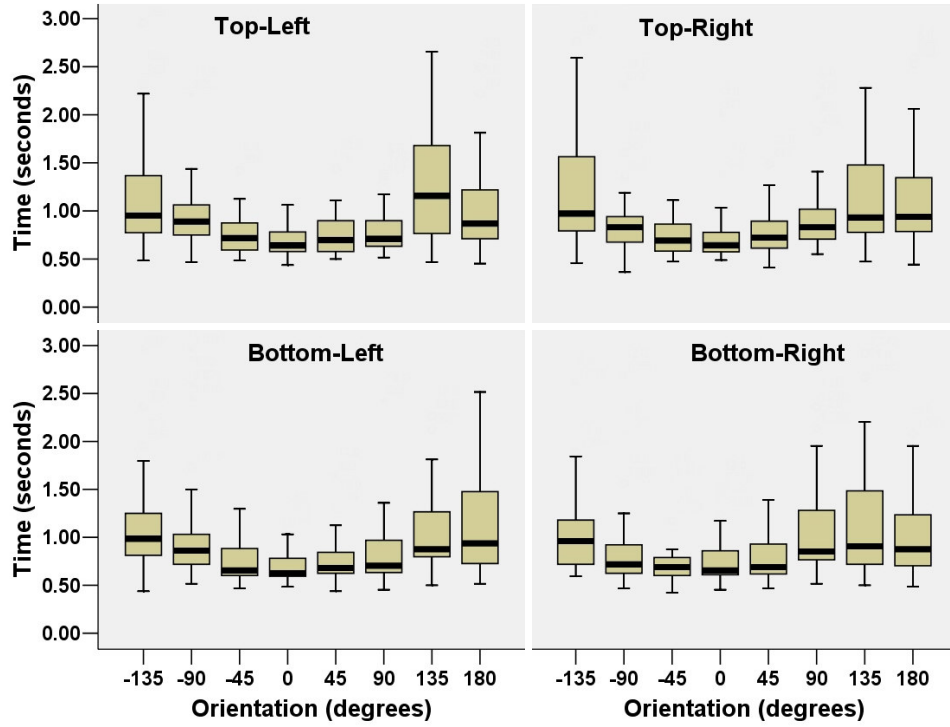


Figure 4. Boxplots of time, in seconds, required to read a single word at each position and orientation. Outliers (>1.5 * IQR) removed.

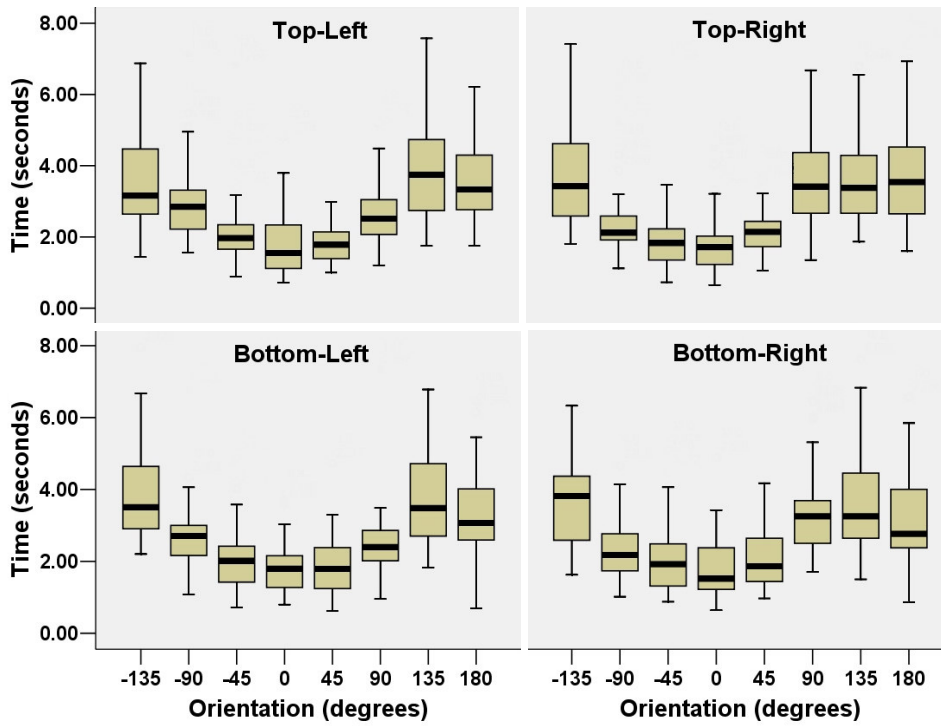


Figure 5. Boxplots of time, in seconds, required to read a 5-6 word phrase at each position and orientation. Outliers (>1.5 * IQR) removed.




Figure 6. Boxplots of time, in seconds, required to read a 6-digit number at each position and orientation. Outliers ($>1.5 * IQR$) removed.

Experiment 2

Goals

A common task in many applications is the serial search: examining many on-screen elements in search of a desired target. Given the results of the previous experiment, we wished to examine the effect of orientation of target and distracters on efficiency in conducting a serial search. We hypothesized that although the amount of rotation of the target and distracters would have a significant effect on performance, the degree of this effect would be less than previous experimental results might suggest. We also hypothesized that a field with more greatly rotated targets and distracters might provide more visual cues, and thus aid in learning. We believed that, over time, this learning would mean that search times for the all-orientations condition would be reduced more than those of the no-orientations condition.

Apparatus

The apparatus used was identical to that in Experiment 1.

Participants

Nine participants recruited from the university community volunteered for the experiment. 7 were male, 2 were female. 6 spoke English natively, while the remaining 3 reported to be excellent readers. Participants received no compensation for their participation.

Procedure

Participants were presented with a search field of randomly positioned 5-6 letter words, each suffixed with a colon and an Arabic numeral between 1 and 3 (eg: “silly : 2”). For each trial, the position, rotation, and text of each of the words remained the same, but the suffixing numeral was randomly changed. For each trial, participants were told to search the field for a given word (different for each trial), and indicate when they had found it by entering its suffixing numeral using the keyboard.

Three datasets were used, which varied only in the degree of orientation of the constituent strings. In the first dataset, each of the strings was presented “right side up”. In the second, strings were drawn at randomly assigned rotations of -45° , 0° , and 45° . In the third, strings were randomly assigned rotations of all 8 multiples of 45° . Figures 7-9 demonstrate the three treatments:



Figure 7. Serial search task with all targets oriented at 0° .

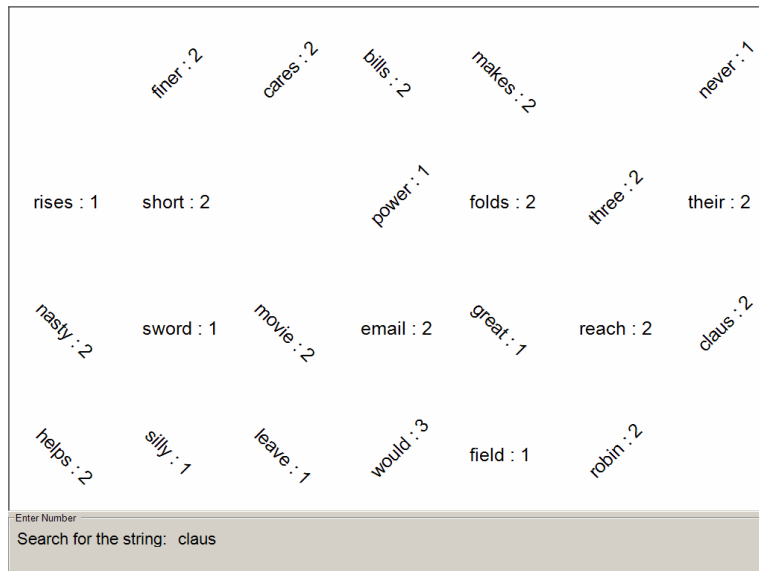


Figure 8. Serial search task with targets oriented at -45° , 0° , and 45° .




Figure 9. Serial search task with targets at all 8 orientations.

Initial instructions were given to the participant as follows:

This experiment is broken into three parts. For each part, you will, 72 times, be presented with a field of strings, each ending with a ":" followed by a number. Each of the 72 times the field will be the same, except that the numbers at the ends of the strings may change. Each time, you will be asked to find a particular string, and indicate to the computer which number it ends with. If you give an incorrect reply, the system will ask you to try again.

Please relax between presentations, but once you have started a trial, find the string and give the input as quickly as you can.

You will be given the opportunity to rest between each presentation. Please do not take long breaks while you are completing a block. You may rest as long as you like between blocks.

The string to find was presented first, and participants then pressed a keyboard button to display the field. Timing began when the field was displayed – Figure 10 demonstrates this sequence:

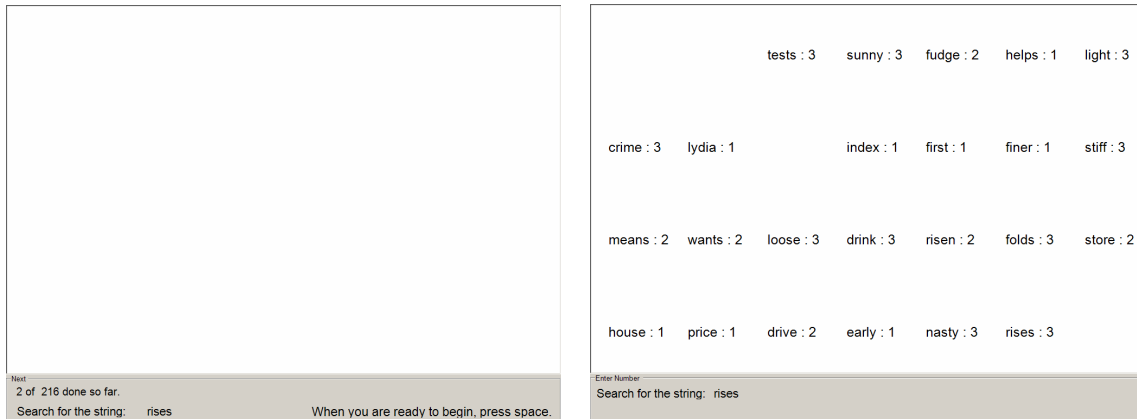


Figure 10. Left: Screen displaying search target to the participant. Right: the search field presented once the subject presses the “space” key, and timing starts for the search.

Timing began when the string was first displayed, and stopped when the subject entered a numeral. Participants were required to enter perfect input, and were told to correct their entry if it was incorrect. Trials with erroneous entries that were subsequently corrected were not included in the analysis.

To ensure that participants understood what was required, they were allowed a practice dataset consisting of several typical searches. They were instructed to conduct as many as was required to become familiar with the apparatus.

Design

Searches were conducted for each of the three datasets. For each dataset, an identical field of 24 5-6 letter words was repeatedly presented. Participants were asked to search for each string in the field on three different occasions within each dataset, resulting in $3 \times 24 = 72$ searches per dataset.

The order of presentation of the datasets was counter-balanced between participants using a Latin-squares design. The assignment of words to treatment, the order of searches, and the makeup (position and orientation) of the field was randomized between participants. In summary, the design was as follows:

- 3 datasets (no rotation, small rotation, complete rotation) X
- 24 strings per dataset (position and orientation randomized) X
- 3 searches per string (search order randomized) X
- 9 participants
- = 1944 searches conducted in total.

Results

As before, trials with erroneous input were discarded from the analysis. There was no significant effect for treatment on error rate ($F_{2,215} = 0.28$, $p = .75$), indicating that orientation did not mislead the subjects. As expected, the orientation of the target and distracter words had a statistically significant effect on the search time for the target ($F_{2,215} = 9.80$, $p < .0001$), though pairwise means comparisons revealed that the search time for the zero and some rotation treatments were not significantly different from one another. As we hypothesized, the difference in performance time between the all-rotations treatment and the others was less dramatic than might have been expected: the mean performance times for the search tasks were 3.3, 3.4, and 3.9 seconds for the no-rotation, small rotation, and complete rotation tasks respectively. Figure 11 is the boxplot for search time for each of the three search tasks.




Figure 11. Boxplot of time required for the serial search under each of the three target / distracter orientation conditions: 0: all oriented towards the user, 1: all oriented at -45° , 0° , or 45° , 2: all oriented at one of the 8-compass positions. Outliers ($>1.5 * IQR$) removed.

Learning of the field did take place: search number had a significant effect on the time required to conduct it ($F_{71,215} = 1.57$, $p < .005$). Contrary to our hypothesis, however, there was no interaction effect between search number and dataset ($F_{142,215} = 7.19$, $p = .84$), indicating that the degree of orientation did not affect the learning of the field.

Discussion

The results of our first experiment confirmed our primary hypothesis: that although significant, the effects of orientation on reading speed are not as large as previous work might have suggested. Where Tinker's (1972) test found a performance penalty of over 200% for reading of a paragraph once text is oriented at 90° , we found penalties of only 26%, 54%, and 17% for short words, phrases, and numbers respectively. Although our results do not refute the notion of a performance cliff at 60° as was found by Koriat and Norman (1985), what is clear is that if such a cliff exists, the depth of the plunge is not nearly as dramatic as they reported.

We attribute the differences in our results to two key experimental differences: the experimental condition, and the experimental task. Unlike the previous work, our participants were free to move their bodies and orient their heads to aid in rotated reading as they would be able to do in real use of tabletop groupware applications. Furthermore, our task required the reading and understanding of words and phrases that were consistently and reliably “real”, so participants could trust their first-glance understanding of the text, rather than second-guess and scrutinize as was required in the Tinker (1972) and Koriat and Norman (1985) experiments.

We were surprised by the finding that 6-digit numbers suffered less from rotation-speed effect than did words and phrases. In our post-experimental interviews, several subjects reported that, when reading numbers, they would begin to type almost right away, and rely on their visual memory of the stimulus as they continued to enter the remaining digits, rather than carefully scanning each digit. We were unable to find references to this behaviour in previous work, which suggests an area of future research.

In the second experiment, our primary hypothesis that orientation would have a significant but minor effect on serial-search, was confirmed. We found that the average search time to find a target word among 23 distracting words suffered only a (statistically insignificant) 3% increase between the zero and some-rotation conditions, and only an 18% increase between zero and all-rotations conditions.

The effect of orientation on task performance is even less significant in the serial search task than it is on the reading task. If we consider the results of single-word reading from the first experiment, and weight the difference in speed of reading from the zero-rotation condition by the proportion of strings in the search field that were at that orientation, we find that a speed difference of 5% in the -45° to 45° condition ($1/3 \times 0\%$ (0°) + $1/3 \times 7.98\%$ (-45°) + $1/3 \times 5.93\%$ (45°) = 4.6%), and a speed difference of 34% ($1/8 \times 64.7\%$ (-135°) + $1/8 \times 26.6\%$ (-90°) + ... $1/8 \times 53.7\%$ (180°) = 33.89%) in the all-rotations condition, would be expected, versus 3% and 18% determined empirically. Thus, serial search of a field of short words is affected less than is the reading of an individual word rotated by the same amount. This finding might be of significance to those conducting general research in the serial-search task.

Although the hypothesized learning did take place as participants continued to make 72 searches within the same field, there was no difference in the learning effect due to text orientation. This was contrary to our secondary hypothesis, which was that when strings were presented in varying orientations, learning of the locations of those strings would be enhanced over a field of uniform rotation. With only three repetitive searches of the same item within a given field and all done within a short timeframe, our experiment was able to only evaluate immediate learning and not long-term recall. It is possible that varying orientations may aid in longer term recall of items, and is worthy of further research.

Conclusions

In collaborative tabletop groupware systems there is a tension between orienting textual data towards the user to facilitate readability, and orientation of interface objects to facilitate interaction, cognition, and communication as observed by Tang (1991) and Kruger et al. (2003). Our experimental results agree with those seen previously that orientation has an effect on speed of reading, and that if maximum reading speed is desired and the position of the intended reader is known, text should be oriented directly towards that user. What we have seen in our study, however, is that if circumstances require it, presenting text at non-optimal orientation does not as severely impair reading performance as previous studies had suggested.

Because the studies by Tang (1991) and Kruger et al. (2003) related primarily to task artifacts in a real-world setting, they offer little guidance as to how system artifacts, such as menus, alerts, and dialogs, should be presented. Our experimental results clearly demonstrate that the effect of orientation on reading labels, numeric and textual data, and performing serial searches, is less dramatic than might have been previously assumed. This indicates that tabletop interface designers could err towards supporting dynamic orientation of artifacts in favour of orienting for readability when they are forced to make a tradeoff between the two choices. Further to this, when presenting shared textual system artifacts, orientation alone does not necessitate that multiple copies or rotating labels be used to allow use by all users, thus saving screen space.

Another way of thinking about this issue is that designers should not attempt to create interfaces that orient text for maximal readability at all costs, or spend all their efforts on such designs, since the impact of slightly rotated text is not as significant as they might naively assume.

Aside from the practical design implications, our work also contributes new data to the empirical literature on readability of oriented text. In contrast to previous work by Tinker (1972) and Koriat and Norman (1984, 1985) that studied the impact of orientation on readability in the task of identifying non-conforming or gibberish strings within other strings, our work provides empirical data for four different tasks: searching for a word in a field of words and reading of non-gibberish words, phrases, and numbers at various orientations. Finally, while the previous research artificially constrained participant's head and body movements, our experiments were carried out in a more natural unconstrained setting, thus providing more ecologically valid data that can be reliably interpreted by designers and experimentalists alike.

Acknowledgments

The authors would like to thank Jonathan Deber, John Hancock, members of the DGP Lab, Mitsubishi Electric Research Labs, experimental participants, and the ECSCW reviewers.

References

- Agrawala, M., Beers, A., McDowall, I., Frohlich, B., Bolas, M. and Hanrahan, P. (1997): 'The two-user Responsive Workbench: support for collaboration through individual views of a shared space', *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co. 1997, pp: 327-332.
- Bruijn, O. and Spence, R. (2001): 'Serendipity within a Ubiquitous Computing Environment: A Case for Opportunistic Browsing', *Proceedings of the 3rd international conference on Ubiquitous Computing*. 2001, pp: 362-370.
- Chapman, J. (1923): *Chapman-Cook speed of reading test*. Ames, IA, Iowa State University Press.
- Dietz, P. and Leigh, D. (2001): 'DiamondTouch: a multi-user touch technology', *Proceedings of the 14th annual ACM symposium on User interface software and technology* Orlando, Florida ACM Press. 2001, pp: 219-226.
- Fitzmaurice, G., Balakrishnan, R., Kurtenbach, G. and Buxton, B. (1999): 'An exploration into supporting artwork orientation in the user interface', *ACM CHI 1999 Conference on Human Factors in Computing Systems*, New York, NY, ACM. 1999, pp: 167-174.
- Huey, E. (1898): 'Preliminary Experiments in the Physiology and Psychology of Reading'. *The American Journal of Psychology*, vol. 9, 1898, pp: 575-586.
- Koriat, A. and Norman, J. (1984): 'What is rotated in mental rotation?', *Journal of Experimental Psychology: Learning, Memory, and Cognition*, vol. 10., 1984 pp: 421-434.
- Koriat, A. and Norman, J. (1985): 'Reading Rotated Words', *Journal of Experimental Psychology; Human Perception and Performance*, vol. 11, number 4, 1985, pp: 490-508.
- Kruger, R., Carpendale, S., Scott, S. and Greenberg, S. (2003): 'How people use orientation on tables: comprehension, coordination and communication', *Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work*. Sanibel Island, Florida, USA, ACM Press, 2003, pp: 369-378.
- Luckiesh, M. (1944): *Light, vision and seeing: a simplified presentation of their relationships and their importance in human efficiency and welfare*. New York, New York, D. Van Nostrand Company Inc.
- MacKenzie, S. and Soukoreff, W. (2003): 'Phrase sets for evaluating text entry techniques', *Extended Abstracts of the ACM CHI Conference on Human Factors in Computing Systems*, 2003, pp: 754-755.
- Matsushita, M., Iida, M., Ohguro, T., Shirai, Y., Kakehi, Y. and Naemura, T. (2004): 'Lumisight table: a face-to-face collaboration support system that optimizes direction of projected information to each stakeholder', *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, Chicago, Illinois, USA ACM Press, 2004, pp: 274-283.
- Ringel, M., Ryall, M., Shen, C., Forlines, C. and Vernier, F. (2004): 'Release, relocate, reorient, resize: fluid techniques for document sharing on multi-user interactive tables', *Extended abstracts of the 2004 conference on Human factors and computing systems*, Vienna, Austria ACM Press. 2004, pp: 1441-1444.
- Rogers, Y., Hazlewood, W., Blevis, E. and Lim, Y. (2004): 'Finger talk: collaborative decision-making using talk and fingertip interaction around a tabletop display', *Extended abstracts of the 2004 conference on Human factors and computing systems*, Vienna, Austria ACM Press, 2004, pp: 1271-1274.
- Rekimoto, J. and Saitoh, M. (1999): 'Augmented surfaces: A spatially continuous work space for hybrid computing environments', *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM Press, 1999 pp: 378-385.
- Shen, C., Lesh, N. and Vernier, F. (2003): 'Personal digital historian: story sharing around the table', *Interactions*, vol. 10, number 2, 2003, pp: 15-22.
- Shen, C., Vernier, F., Forlines, C. and Ringel, M. (2004): 'DiamondSpin: an extensible toolkit for around-the-table interaction', *Proceedings of the SIGCHI conference on Human factors in computing systems*, 2004, pp: 167-174.

- Streitz, N., Prante, T., Muller-Tomfelde, C., Tandler, P. and Magerkurth, C. (2002): 'Roomware©: the second generation CHI '02 extended abstracts on Human factors in computing systems', Minneapolis, Minnesota, USA ACM Press, 2002, pp: 506-507.
- Streitz, N., Geuler, J., Holmer, T., Konomi, S., Miller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P. and Steinmetz, R. (1999): 'i-LAND: an interactive landscape for creativity and innovation', *Proceedings of the SIGCHI conference on Human factors in computing systems: the CHI is the limit*, Pittsburgh, Pennsylvania, United States ACM Press, 1999, pp: 120-127
- Stroop, J. (1935): 'Studies of interference in serial verbal reactions', *Journal of Experimental Psychology: General*, vol. 18, pp: 643-662.
- Tang, J. (1991): 'Findings from observational studies of collaborative work', *International Journal of Man-Machine Studies*, vol. 34, number 2, 1991, pp: 143-160.
- Tinker, M. (1972): 'Effect of angular alignment upon readability of print', *Journal of Educational Psychology*, vol. 47, 1972, pp: 358-363.
- Wu, M. and R. Balakrishnan (2003): 'Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays', *Proceedings of the 16th annual ACM symposium on User interface software and technology*, Vancouver, Canada ACM Press, 2003 pp: 193-202.