Cognitive Engineering Laboratory

Evaluation of a Display Design Space: Transparent Layered User Interfaces

Beverly L. Harrison, Hiroshi Ishii, Kim J. Vicente, and Bill Buxton

CEL 94-07

Cognitive Engineering Laboratory University of Toronto Department of Industrial Engineering
4 Taddle Creek Rd. Toronto, Ontario, Canada M5S1A4
phone: (416) 978-0881 email: benfica@ie.utoronto.ca fax: (416) 978-3453

Cognitive
Engineering
Laboratory

COGNITIVE ENGINEERING LABORATORY

Director: Kim J. Vicente, B.A.Sc., M.S., Ph.D.

The Cognitive Engineering Laboratory (CEL) at the University of Toronto (U of T) is located in the Department of Industrial Engineering, and is one of three laboratories that comprise the U of T Human Factors Research Group. The CEL began in 1992 and is primarily concerned with conducting basic and applied research on how to introduce information technology into complex work environments, with a particular emphasis on power plant control rooms. Professor Vicente's areas of expertise include advanced interface design principles, the study of expertise, and cognitive work analysis. Thus, the general mission of the CEL is to conduct principled investigations of the impact of information technology on human work so as to develop research findings that are both relevant and useful to industries in which such issues arise.

Current CEL Research Topics

The CEL has been funded by Atomic Energy of Canada, Ltd., Natural Sciences and Engineering Research Council, Defense and Civil Institute for Environmental Medicine, Japan Atomic Energy Research Institute, and Asea Brown Boveri Corporate Research - Heidelberg. The CEL also has collaborations and close contacts with the Westinghouse Electric Company, and Toshiba Nuclear Energy Laboratory. Current projects include:

- Studying the interaction between interface design and skill acquisition in process control systems.
- Understanding control strategy differences between people of various levels of expertise within the context of process control systems.
- Developing a better understanding of the design process so that human factors guidance can be presented in a way that will be effectively used by designers.
- Evaluating existing human factors handbooks.
- Developing advanced interfaces for computer-based anesthesiology equipment.

CEL Technical Reports

For more information about CEL, CEL technical reports, or graduate school at the University of Toronto, please contact the address printed on the front of this technical report, or send email to Dr. Kim J. Vicente at <benfica@ie.utoronto.ca>.

Evaluation of a Display Design Space: Transparent Layered User Interfaces

Beverly L. Harrison[1] *Hiroshi Ishii*[2] *Kim J. Vicente*[1] *Bill Buxton*[3]

[1]Dept of Industrial Engineering
University of Toronto
Toronto, Ontario, Canada
M5S 1A4
(416)978-1531
beverly@dgp.utoronto.ca
benfica@ie.utoronto.ca

[2]NTT Human Interface lab
1-2356 Take, Yokosuka-Shi
Kanagawa, 238-03 Japan
ishii@chi.xerox.com

[3]Alias Research Ltd., and
Dept. of Computer Science
University of Toronto
Toronto, Ontario, Canada
M5S 1A4
buxton@dgp.utoronto.ca

ABSTRACT

This paper proposes a design space for displays which varies either the degree of transparency/opaqueness of layers, or the depth between layers (in planes), or a combination of both. We first describe the dimensions of our design space and relate it to prior work. We then describe two experiments run to evaluate transparency, task interference, and legibility. The results of our experiments suggest design parameters which we are using in several real-world applications. Some insightful comments from the users of these applications are summarized. Finally, we outline our future research plans.

KEYWORDS: display design, empirical evaluation, transparency, user interface design, interaction technology

INTRODUCTION

Limited screen real estate combined with graphical interface design has resulted in systems which show a proliferation of overlapping windows, menus, dialog boxes, and tool palettes. It is not feasible to "tile" our computer workspaces to facilitate keeping track of things. There are too many objects.

One way to work within existing space constraints, perhaps with less loss of visual awareness, is to design systems where transparent objects are possible. If we could see the "top-most" object while simultaneously seeing *through it* to underlying or overlapped objects, we might preserve more sense of context and content. Changing between the layers would merely require a shift of visual attention.

Several key design issues need to be investigated if users are expected to visually attend to two superimposed images. Can users selectively attend to a chosen "layer" without visual interference from the other? Are there certain display characteristics or task properties which facilitate or preclude overlapping displays?

This paper proposes a design space and summarizes empirical evaluation of this space. In this design space, layers are differentiated by either the degree of transparency/opaqueness, or by the perceived planar depth between layers (not continuous 3-D), or by both depth and transparency level. The experiments described here investigate the levels of task interference using varying degrees of transparency to partially address some of the design issues. We then summarize user feedback from several real-world applications which we have built using transparent interfaces.

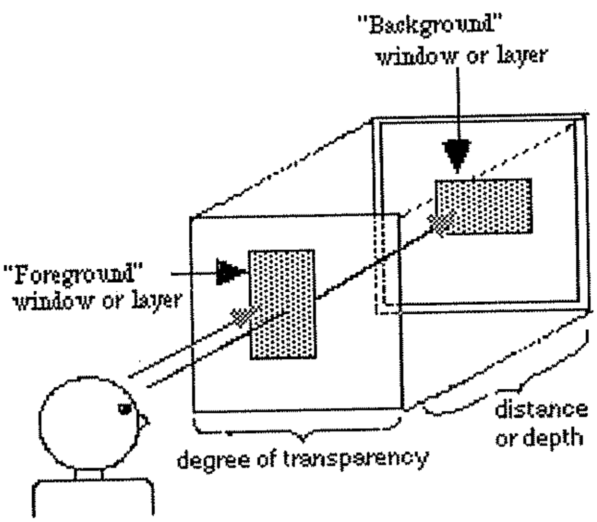

Figure 1. Concept of Layered Displays

TRANSPARENT LAYERED DISPLAY DESIGN SPACE

We propose a design space for managing two or more sources of visually presented information simultaneously (Figure 1 and Figure 2).

In one dimension, we vary the level of transparency/opaqueness between the two displays. *Fully opaque objects* reflect traditional window, palette, and menu design in current graphical user interfaces. *Fully transparent designs* reflect some of the more advanced interfaces such as those used in Heads Up Displays (HUDs) in aviation or in the Clearboard system (Ishii and Kobayashi, 1991). In HUD design, aircraft instrumentation (a graphical computer interface) is superimposed on the external real world scene, using specially engineered windshields. In the Clearboard work, a large drawing surface is overlayed on a video image of the user's collaborative partner. *Semi-transparent designs* include such things as video overlays (like those used in presenting sports scores while the game is playing), "3-D silk cursors" (Zhai et al, 1994) or Toolglass–like tool palettes (Bier et al, 1993; Kabbash et al, 1994).

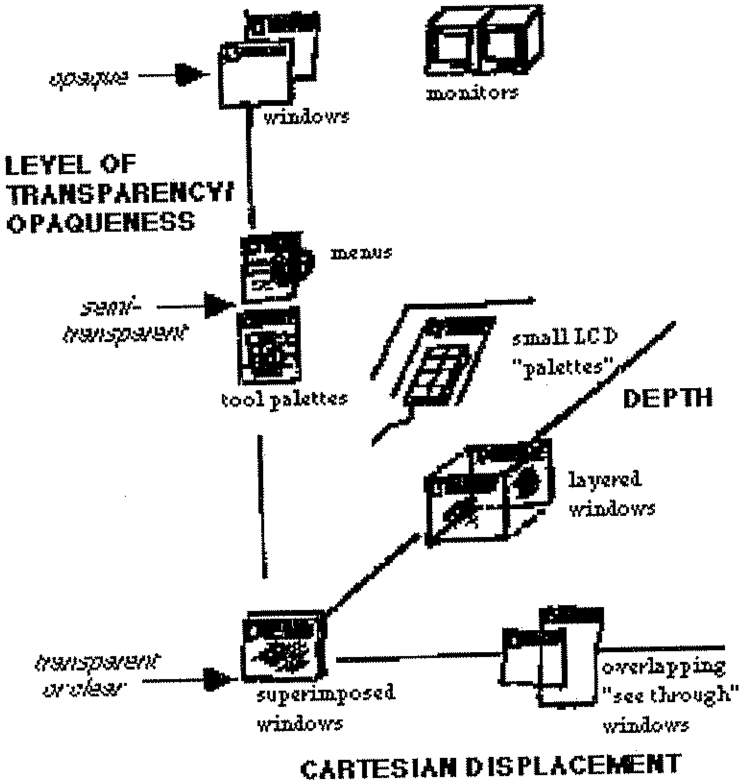

Figure 2. Design Space Dimensions

Along another dimension we vary the perceived depth of the planes between two displays, where one image appears closer to the user while the other is in the background. This is accomplished using half-silvered mirrors, polarizing filters, or special transparent LCD displays. In this case the user looks through the display presented in the foreground to see the display presented in the background (Kobayashi and Ishii, 1994). There are limited examples of such systems. Knowlton (1976) used graphical overlays projected downwards onto half-silvered mirrors over blank keyboards to dynamically re-label buttons and functions keys (e.g., for telephone operators). Schmandt (1983) built a system to allow users to manually manipulate and interact with objects in a 3-D computer space using a 3-D wand. Again a half–silvered mirror was used to project the computer space over the user's hand(s) and the input device. Disney has also come out with a product called the "ImaginEasel" for animators and artists which keeps the user's hand and input device in the workspace (using mirrors). Layers in this design space are distinguished by either transparency or depth or both.

To summarize, the small amount of display real estate available relative to the amount of data to be displayed presents a real challenge to user interface design. To date, two main strategies have been applied to the problem. In the first, the screen is partitioned, or *tiled,* into a number of non-overlapping windows. This we refer to as the *space multiplexed* strategy. In the second, windows lie on top of one another. Only the top one is visible at any given time, but a mechanism is provided to rapidly change which window is visible. This we refer to as the *time multiplexed* strategy. Most frequently, a hybrid of the two is used. What we propose in this paper, however, is a third strategy. Through the use of transparency in the background of windows, the contents of windows underneath others is visible, or at least partially visible. This "new" strategy we refer to as *depth multiplexing.*

On the one hand, the depth multiplexing approach offers the best of both worlds: windows need not be tiled to be visible. Hence, ideally, less information is obscured. On the other hand, the potential for content of one window interfering with another above or below it is introduced. Our prototypes show clearly that in some situations the technique works well, while in others there are real problems. The objective of our research agenda, of which the current paper is a part, is to determined a more formal understanding of the constraints of such an approach.

WHERE MIGHT TRANSPARENT LAYERED DISPLAYS BE USED?

In most graphical user interfaces (GUIs), a number of situations arise as part of our day-to-day work which require us to either keep track of several things at once (in parallel) or to stay aware of an interrupted or obscured task while dealing with a new task (in serial). A number of such situations are outlined below, reflecting the diverse range of possible applications.

- working on a document when we get interrupted by a dialog box or warning message
- a pull-down menu (or pie menu) may temporarily block part of our current window. (The selected menus items may go on to create further dialog boxes of their own.)
- using tear-off tool palettes (which behave as tiny overlapping windows)
- collaborating with a partner shown in a video window and we want to glance at both our partner and the work in progress (e.g., a drawing or document)
- viewing a live video conversation with one person while monitoring several Portholes-like connections to others for peripheral awareness of their availability
- using a help system which displays a window of information and we would like to remember the context from which we requested help
- looking at an overview diagram of a system and a more detailed view simultaneously, or a view of two different parts of a graphical representation of a system
- using an interactive dialog box to change the drawing or modeling characteristics of an underlying image, model, or animation

These situations all share a common attentional problem: we need to be visually aware of multiple objects which overlap and obscure each other. All of these scenarios have two (or more) "tasks". In some cases we either wish to time-share the two tasks or selectively attend to one task excluding the other, with minimal switching effort.

By their very nature, many of the proposed task pairs have an implicit active and passive task, or foreground and background task. We need a *peripheral awareness* of the passive or background task while we temporarily divert most of our attention to the active or foreground task.

WHY DO WE BELIEVE THAT THIS MIGHT WORK?

A number of areas have contributed to our understanding of selective and divided attention, selective looking, and display design. The research findings reported thus far in other domains suggest that this new design space is a promising method of presenting multiple sources of visual information.

Kohler (1972) originally investigated selective looking (monitoring dual tasks) by building headgear using half-silvered mirrors which presented the scene of the world in front of him superimposed on the scene of the world behind him. He reported that he could easily switch between these two views; the unattended scene seemed to "disappear" from sight.

Motivated by this work, further studies were carried out (Neisser and Becklen, 1975; Becklen and Cervone, 1983) using two overlayed or superimposed video images presented on a single monitor. In the first study (Neisser and Becklen, 1975) the tasks were visually distinctive: a hand slapping game and a ball tossing game. In the later study (Becklen and Cervone, 1983) both tasks were visually similar ball tossing games; the tasks were differentiated by the color of the shirts worn by the players. In both cases, subjects were asked to monitor one task and indicate the irregular occurrence of target events in this task. Meanwhile, bizarre events were sporadically presented in the non-monitored task.

Subjects were easily able to monitor the target task to the exclusion of the unattended task. Subjects did not notice the bizarre events, even when the experiment was stopped during or immediately after the bizarre event occurred and the subjects were asked about it. This result still held when the bizarre event was presented in the exact same visual location where the target event occurred (i.e., within foveal range). This seems to indicate that the intentionally unobserved task goes virtually unnoticed. A number of alternative explanations for this phenomenon were discussed and discounted. This work suggests that two superimposed *video* tasks can be easily monitored with minimal interference. However, the extent of simultaneous task awareness is unclear.

Similar results in selective looking have been found in studies of dual task monitoring in HUDs typically used in aircraft control and navigation tasks. Specific advantages cited include improved flight performance, superior object tracking, (Wickens and Larish, 1991; Long and Wickens, 1994). The primary disadvantage is "attentional tunneling" – fixation on the HUD to the exclusion of events in the real world, particularly unexpected events (or vise-versa) (Wickens et al, 1993). Again subjects are easily able to differentiate either display layer easily. Practice seems to improve *simultaneous* monitoring performance.

This previous research, though not applied directly to graphical user interface design per se, suggests promising evidence for the use of superimposed transparent displays. Based on these results, one would anticipate reduced switching time and improved awareness by minimizing head and eye movement and re-focusing. Also one can reasonably anticipate that users will be able to treat the sources separately and voluntarily attend to one or the other (with varying degrees of interference). Clearly the usability of the resulting designs will also depend upon the information or task characteristics.

As in most interface designs, one can anticipate some inappropriate applications and pitfalls as well. In cases where "missed observations" have a high cost, reducing visibility through transparency might be undesirable. Also if both tasks must be simultaneously monitored and both have high attentional demands, the attentional tunneling problems might arise. Finally, while this would seem feasible for distinctive types of information, we must evaluate how well this technique works for visually similar information types.

EXPERIMENT — TRANSPARENCY EFFECTS ON TASK INTERFERENCE AND LEGIBILITY

One of the strongest, most consistent measures of task interference is the Stroop Effect (Stroop, 1935). It is virtually impossible to consciously block or prevent the Stroop Effect in selective looking tasks, despite numerous experimental permutations (over 700 articles to date – for reviews see Jensen and Rohwer, 1996; MacLeod, 1991). In traditional Stroop tasks, a series of random words are presented which appear in a randomly chosen color of ink (e.g., red, green, blue, yellow). Subjects must name the *ink color* while ignoring the word. Some of the words are neutral (e.g., uncle, shoe, cute, nail); other words are the names of conflicting colors (e.g., yellow, blue, green, red). Consistent, significant performance degradation occurs when conflicting color words are used and subjects attempt to name the color of the ink (e.g., the word "red" appears in green ink; the correct response is green). In later studies (e.g., Kahneman and Chajczyk, 1983), a consistent and significant Stroop Effect was found even when the word was printed in black ink, presented adjacent to a color bar.

Our experiments test how varying *transparency* effects interference between the displayed word and the color target, using a traditional Stroop test (Figure 3). The word is seen by looking "through" the color patch. The degree of transparency of the color patch is varied from clear (high interference – both items are completely visible) to opaque (no interference – the word is invisible, reflecting the baseline condition of naming the color when only the color is present). The color naming experiment tests the extent of interference from the word and how transparency levels might dilute this interference. The word naming experiment evaluates performance and text legibility. When combined, the two experiments suggest interface design parameters where interference is minimized and the word is still fairly legible (awareness is preserved).

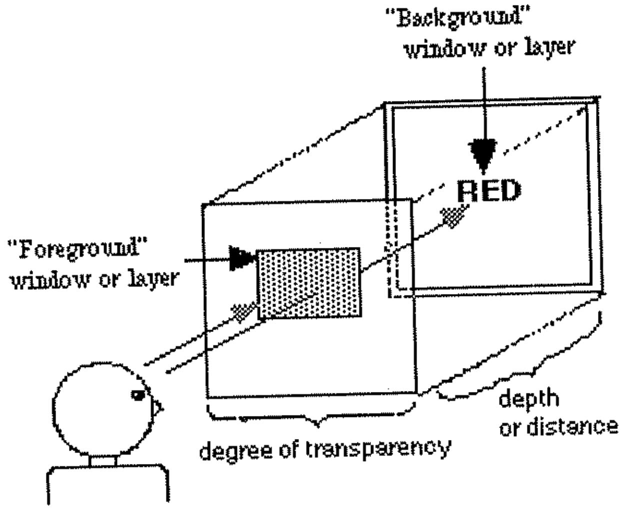

Figure 3. Layered display design[1]

Hypotheses

H1: As transparency level is decreased (i.e., seeing the word through the color patch gets more difficult) the response time is decreased (less interference and therefore improved performance), in the color naming experiment.

H2: At some transparency level > 0% (where the word is still visible) performance will be equivalent to transparency = 0%, in the color naming experiment.

H3: Word naming performance will be inversely related to transparency level (i.e., as transparency increases the word gets harder to see and therefore takes longer to read).

Experimental Design: Color Naming Experiment and Word Naming Experiment

For our experiments we used 4 colors: red, blue, green, and yellow. Words (helvetica, 78 point, uppercase) appeared "through" the colored rectangular patch. We used neutral words UNCLE, NAIL, CUTE, and FOOD in addition to the four color names. Transparency levels were varied as: 0% (baseline condition - only one of the word or color shows), 5%, 10%, 20%, 50%, 100% (clear - both the word and color show). Images

---

1 Figures are modified versions of those which originally appeared in Kobayashi and Ishii, 1994.

were produced using alpha blending on an SGI Indigo™ to ensure accurate transparency levels.

Both experiments used the same methodology and image files; however, one experiment was a color naming task and the other was a word naming task. Order of presentation was counter–balanced and spaced one day apart. No cross task interference is anticipated (MacLeod and Dunbar, 1988). The word naming experiment baseline condition was a word only – presented with no color patch. The color naming experiment baseline condition was a color patch only – presented with no word. There were no other differences between the two experiments. (The word naming experiment should not have any Stroop effects but performance should be effected by the visibility of the word.)

A fully randomized, within subject, repeated measures design was used. There were 4 conditions: non-conflict or neutral (the word was a neutral word), incongruent color (a conflicting color word was present), congruent color (the color word matched the color of the patch), and baseline (color or word only). The latter baseline condition was taken as a transparency level = 0%. The remaining 5 levels of tranparency (5%, 10%, 20%, 50%, 100%) were applied across the other 3 non-baseline conditions. A total of 180 unique combinations resulted. For each subject, 3 runs of the entire set of 180 images were shown. Trials were presented in random order at 5 second intervals. Subjects used a button press during the inter-trial-interval to indicate an error in response. (Subject logged errors showed no significant difference from experimenter logged errors in previous studies, e.g., MacLeod, and Dunbar 1988. Videotapes of our experiment were later analyzed to confirm this.) Each experiment lasted about 45 minutes.

Experimental System Configuration

The experiments were run using the PsyScope software and hardware (Cohen et al., 1993) with a headset microphone on a Macintosh IIfx. Audio levels were adjusted before each subject was run. Verbal responses were logged within 1 msec of accuracy. Subjects sat at a fixed distance of 100cm from the screen.

Procedure

Subjects were given 20 practice trials to familiarize them with the operation of the experiment. These trials were randomly selected from the set of 180 possible combinations. Following this, subjects were shown 3 sets of the 180 combinations (15 minutes per set), with rest breaks in between each set.

Subjects were debriefed at the end of the experiment. Open ended comments were recorded and the experiment was video and audio taped for analysis purposes. Response times and subject ranked errors were logged by the computer.

Subjects

A total of 16 students from the University of Toronto were run as subjects[1]. They were pre-screened for color-blindness. Subjects were paid for their participation and could voluntarily withdraw without penalty at any time.

RESULTS – COLOR NAMING TASK

A univariate repeated measures ANOVA was run (with adjusted error terms). Table 1 summarizes the statistically significant results:

Source	df	F value	p value <
subj*time	6	2.3	.05
transp	4	4.4	.05
color	3	9.4	.001
transp* color	12	2.05	.05
word	8	5.18	.001
transp* color* word	96	1.3	.05

Table 1. ANOVA - color naming.

Transparency was both a significant main effect and a significant interaction effect when taken with color, or color and word type. This suggests that transparency may indeed dilute the interference/Stroop effect. The mean response times reflect this: 100% - 713 ms, 50% - 687 ms, 20% - 680 ms, 10% - 624 ms, 5% - 624 ms, 0% - 603 ms. Post-hoc analyses were carried out to compare means for the transparency * color and transparency * color * word interactions (Student-Newman-Keuls test with alpha levels = .05).

For levels of transparency of 5% and 10% (word was only slightly visible), the means across all word types

1. This paper reports the results of 4 subjects run thus far. Statistical significance is expected to improve further upon completion of the analysis.

are statistically equivalent (no interference/Stroop effect). This suggests that levels of 5% and 10% have minimal or no task interference. At levels above 20%, three groupings of means occur (as the Stroop effect would predict): blank and congruent, neutral words, and incongruent words. As transparency levels are further increased these differences become more pronounced.

Similarly post-hoc analysis of color indicates that the hue and saturation of the color may effect interference slightly (likely by effecting legibility of the word in some subtle way). This effect showed up at the 20% and partially at the 50% transparency levels only. Interference effects seemed strongest for yellow, followed by red/green, then blue. For higher levels all colors produced similar results.

Subject errors in response occurred only occasionally (average of 12 per 540 trials) and almost exclusively on the color-incongruent trials (average 10 incongruent, 2 neutral). Error trials were not used in the above analysis.

RESULTS – WORD NAMING TASK

A univariate repeated measures ANOVA was run (with adjusted error terms). Table 2 summarizes the statistically significant results.

Once again transparency was both a significant main effect and a significant interaction effect when taken with color, or color and word type. This suggests (not surprisingly) that transparency effects word naming performance. Post-hoc analyses were carried out to compare means for the transparency * color and transparency * color * word interactions (Student-Newman-Keuls test with alpha levels = .05).

Source	df	F value	p value <
subj*time	6	271.6	.001
transp	4	4.32	.05
color	3	207.2	.001
transp* color	12	2.37	.05
transp* color* word	96	2.37	.001

Table 2. ANOVA - word naming.

For levels of transparency of 5% subjects reported great difficulty in seeing the word, so few trials were recorded. (Subjects reported "none" when they could not make out the word.) At these levels, certain colors performed better than others. Yellow and green hues seemed "easier" to detect the word on (as shown in the post-hoc analysis of means). At 10%, about 1/3 of the trials were not properly visible to the subjects. Color made no difference and performance was equivalent across all word types (as predicted by Stroop, 1935 and others). At 20% levels and higher, all words were easily read.

Subject errors in response occurred only rarely (average of 2 per 540 trials) and almost exclusively on the color-incongruent trials. Error trials were not used in the above analysis.

DISCUSSION

We have found that degree of transparency dilutes the interference/Stroop effect in a monotonic fashion until levels < 20% are used. At levels of 5% and 10% minimal or no interference seems to take place (using means comparison tests). As the levels increase beyond 20% the interference effects are more pronounced (the differences between the means is greater). This seems to confirm our first 2 hypotheses.

Word naming seems highly error prone at levels of 5%. At levels of 10% subjects could accurately name most of the words, though they seemed to perform slightly better depending upon what the background color was. It seems that there must be some interactions between hue/saturation and legibility – despite using the same alpha blending technique to create exact transparency levels across all colors. This suggests that certain colors might be more profitably used in transparent windows or interfaces - though this remains to be tested. In general, word naming performance supports hypothesis 3 (above) and is impervious to the Stroop effect (as expected).

USER FEEDBACK FROM REAL APPLICATIONS

We installed transparency into some interactive dialog boxes within a 3-D modeling/animation system. In this system the user needs to see a potentially large model (full screen) while changing various attributes of the model or of the drawing tools. These changes are done by selecting options in a movable window or dialog box. Typically a user might have 3 or 4 such interactive dialog windows open at all times.

We had several users of varying levels of expertise evaluate the transparent windows. We also asked users to select a "personal favorite" transparency level using a slider bar. Substantial in-depth investigation is still being conducted. However, several insightful comments have already been noted.

The degree of visual distinction between the two tasks strongly influences the extent of possible interference and perceived difficulty. Users found transparent windows (text, buttons) were easier to use over solid models/images than those superimposed over wire frame drawings. Higher levels of "opaqueness" seemed to partially compensate in the more difficult task situation.

As familiarity with the interactive window layout improved, users preferred corresponding increases in transparency. They preferred to see "less" of the window and more of the underlying image. The window items were needed only as outlines to target selections - the actual legibility of the text was substantially less important. This suggests that border of windows and buttons and data entry areas might be handled in a different way than the actual names and labels.

CONCLUSIONS AND FUTURE DIRECTIONS

We have illustrated a method of empirically testing our design space dimensions using well-established theoretical measures. We are currently conducting experiments that evaluate the planar depth dimension of our design space. The Stroop color naming and word naming tasks are being used to evaluate how varying the distance between display planes might effect task interference.

From our experimental work described here, we are now testing user performance on transparency levels between 10% to 20% in a variety of real world tasks using transparency in menus, dialog boxes, and windows.

Drawing on preliminary findings described above, we are evaluating whether the entire overlapping window should be transparent or only the background of the window, leaving the text and/or borders visible. We are also evaluating whether the transparency level should automatically change for the entire window (or a portion of the window) when the cursor is moved over it. In this case, items would become more opaque when they might be more "active" or selectable (dynamic transparency).

We believe that interface designers can take advantage of both the intrinsic properties of the task and of an understanding of human visual attention to design new display techniques and systems. The design space proposed in this paper supports the idea of active/passive or foreground/background tasks by providing users with an awareness of one task while they focus on the other. In this way, inherent characteristics of the task are supported in the interface while providing enhanced functionality. This design applies across numerous situations and applications. We believe that results thus far show promising advantages for creating new user interfaces and interaction techniques. We are exploiting possibilities of new technology in a way that is sensitive to both psychological and task constraints.

ACKNOWLEDGEMENTS

Support for our laboratory is gratefully acknowledged from the Natural Sciences and Engineering Research Council (NSERC), Apple Computer, the Information Technology Research Centre (ITRC), Xerox PARC, Alias Research Inc., and the Ontario Telepresence Project. We would also like to thank Professor Allison Sekuler and the Psychology Dept. for the generous use of their software and equipment, Prof. Colin MacLeod for comments on the Stroop design used here, and members of the Graphics Lab and Congitive Engineering Lab for programming consultation. Programming for the system shown in the accompanying video tape was done by Shumin Zhai.

REFERENCES

Becklen, R. and Cervone, D. (1983) Selective looking and the noticing of unexpected events. *Memory and Cognition*, 11 (6), pp. 601-608.

Bier, E. A., Stone, M. C., Pier, K., Buxton, W., and DeRose, T. D. (1993). Toolglass and magic lenses: The see-through interface. *Proceedings of SIGGRAPH'93*. Anaheim, CA.

Cohen J.D., MacWhinney B., Flatt M. & Provost J. (1993). PsyScope: A newgraphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments & Computers,* 25(2), pp. 257-271.

Ishii, H. and Kobayashi, M. (1991). Clearboard: A seamless medium for shared drawing and conversation with eye contact. *Proceedings of CHI'91,* Monterey, CA, pp. 525-532.

Jensen, A. R. & Rohwer, W. D. (1966). The Stroop color-word test: A review, *Acta Psychologica,* 25, pp. 36-93.

Kabbash, P., Buxton, W. A. S., and Sellen, A. (1994). Two-handed input in a compound task. *Proceedings of CHI'94,* Boston, MA., pp. 417-423.

Kahneman, D. and Chajczyk, D. (1983) Tests of the Automaticity of Reading: Dilution of Stroop Effects by Color-Irrelevant Stimuli. *Journal of Experimental Psychology: Human Perception and Performance,* 9 (4), August, 1993, pp. 497-509.

Keele, S. W. (1972) Attention demands of memory retrieval. *Journal of Experimental Psychology*, 93, pp. 245-248.

Knowlton, K. C. (1977). Computer displays optically superimposed on input devices. *Bell System Technical Journal*, Vol. 56, No. 3, March, 1977. pp. 367-383.

Kobayashi, M. and Ishii, H. (1994). DispLayers: Multi-Layer Display Technique to Enhance Selective Looking of Overlaid Images. Poster from *CHI'94 Conference*, April, 1994, Boston MA.

Kohler, P. A. (1972) *Aspects of Motion Perception.* New York: Pergamon.

Larish, I and Wickens, C. D. (1991). Divided Attention with Superimposed and Separated Imagery: Implications for Head-Up Displays. *University of Illinois Institute of Aviation Technical Report* (ARL-91-4/NASA HUD-91-1).

Long, J. and Wickens, C. D. (1994). Implications of Object vs. Spaced Based Theories of Attention in the Design of the Aircraft head-Up Display. *University of Illinois Institute of Aviation Technical Report* (ARL-94-6/NASA HUD-94-1).

MacLeod, C. M. (1991). Half a Century of Research on the Stroop Effect: An Integrative Review. *Psychological Bulletin,* Vol. 109, No. 2, pp. 163-203.

MacLeod, C. M. and Dunbar, K. (1988). Training and Stroop-like interference: Evidence for a continuum of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, pp.126-135.

Neisser, U. and Becklen, R. (1975) Selective Looking: Attending to Visually Specified Events. *Cognitive Psychology*, 7, pp. 480-494.

Posner, M. I. and Snyder, C. R. R. (1975) Facilitation and inhibition in the processing of signals. In P. M. A. Rabbitt and S. Dornic (Eds.), *Attention and Performance V.* New York: Academic Press.

Schmandt, C. (1983). Spatial input/display correspondance in a stereoscopic computer graphic workstation. *Computer Graphics*, Vol. 17, No. 3., pp.253-259.

Stroop, J. R. (1935). Factors affecting speed in serial verbal reactions. *Journal of Experimental Psychology,* 18, pp. 643-662.

Wickens, C. D., Martin-Emerson, R., and Larish, I. (1993). Attentional tunneling and the Head-up Display. *Proceedings of the 7th Annual Symposium on Aviation Psychology,* Ohio State University, Ohio, pp.865-870.

Zhai, S., Buxton, W., and Milgram, P. (1994). The "silk cursor": Investigating transparency for 3D target acquisition. *Proceedings of CHI'94,* Boston, MA., pp. 459-464.