Cognitive
Engineering
Laboratory

Semi-Transparent "Silk" User Interface Objects: Supporting Focused and Divided Attention

Beverly L. Harrison, Shumin Zhai, Kim J. Vicente, and Bill Buxton

CEL 94-08

VELUT ARBOR ÆVO

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

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>.

Semi-Transparent "Silk" User Interface Objects: Supporting Focused and Divided Attention

Beverly L. Harrison[1] *Shumin Zhai*[2,1] *Kim J. Vicente*[1] *Bill Buxton*[2,3]

[1]Dept. of Industrial Engineering
University of Toronto
Toronto, Ontario, Canada
M5S 1A4
tel: (416)978-1531
fax: (416)978-4765
beverly@dgp.utoronto.ca

[2]Alias Research Inc.
110 Richmond Street East
Toronto, Ontario, Canada
M5C 1P1

[3]Dept. of Computer Science
University of Toronto
Toronto, Ontario, Canada
M5S 1A4

Contact person: Beverly Harrison

ABSTRACT

This paper proposes a framework for designing and evaluating user interfaces with semi-transparent "widgets", including windows, menus, dialogue boxes, screens, or other objects. "See-through" interface objects (or "silk widgets") preserve our sense of awareness and context while minimizing attentional switching costs between layers of objects. However, as in most interface designs, semi-transparency also introduces unique challenges. We identify the trade-off between divided and focused attention as the central issue in designing semi-transparent interfaces. We first describe relevant theoretical concepts in focused and divided attention and a derived methodological framework to assist further investigation. Under this framework, two complimentary approaches, controlled experimental evaluation and observation through prototyping, have been taken to study how semi-transparent widgets can improve or degrade human computer interaction performance. This paper presents our motivation, design framework, evaluative approach, and the preliminary results of our investigation.

KEYWORDS: graphical user interface (GUI), display design, transparency, user interface design, interaction techniques, divided attention, focused attention.

INTRODUCTION

This paper describes aspects of a new research program where we are investigating graphical user interfaces from an attentional perspective (as opposed to a more traditional visual perception approach). The central research issue is how we can better support both focusing attention on a single interface object without distraction from other objects, and supporting dividing or time sharing attention between multiple objects to preserve context or awareness. This attentional trade-off seems to be a central, but as yet comparatively ignored issue, in many interface designs.

Current graphical user interfaces (GUIs) provide us many interactive "widgets" within a relatively small physical space determined by the screen size. These widgets include such things as menus, windows, dialogues, icons, and tool palettes where each widget may have a context or relationship to some of the other objects. These widgets facilitate the dialogue of human computer interaction. They are the key components of modern GUIs. However, the small amount of display real estate available relative to the amount of data to be displayed presents a real challenge to the design and allocation of these widgets.

To date, two main strategies have been applied to the problem. We have classified these strategies as: space multiplexing (widgets are accessed by spatial location) and time multiplexing (widgets are accessed in sequence temporally). We propose a third strategy: depth multiplexing (widgets are transparently overlayed in layers).

In the first, the screen is partitioned, or *tiled*, into a number of non-overlapping windows. The advantage of this approach is that users can clearly see and identify each widget. In order to switch attention from one widget to another, users need only change their focus of attention spatially. This we refer to as the *space multiplexed* strategy. The disadvantage of this strategy is that a limited number of objects can be shown at any given time, leaving many objects temporarily "hidden". These hidden or obscured objects must be explicitly found and brought to the foreground, replacing (in spatial location) some existing item. Also the relationships or context between objects is difficult to represent by simple rules of adjacency.

In the second strategy, widgets lie on top of one another (*overlapped*). Only the top object or widget is visible at any given time, but a mechanism is provided to rapidly change which object is visible. This is typically some combination of mouse clicking, selection, and dragging. At any one moment in time a subset of widgets are visible. This we refer to as the *time multiplexed* strategy. This strategy has the advantage of representing many more objects but at the expense of increased "attentional switching costs" and a decrease in awareness of which objects are present (some are partially or completely occluded). Most frequently, a hybrid of the two strategies is used.

What we propose in this paper, however, is a third strategy. Through the use of semi-transparency in the background of widgets, the contents of the objects underneath are visible, but with lower contrast (Figure 1, colour plate). The effect of semi-transparent background of the widgets resembles the visual characteristics of silk stocking material. We therefore name the semi-transparent interface objects "silk" widgets.

Silk widgets can still be tiled or overlapped. When overlapped the user can see through the top-most object to preserve a sense of awareness (and a sense of context) for the object in the underlying layer. This "new" strategy we refer to as *depth multiplexing.* It can also form a hybrid interface of opaque and transparent objects, combined with tiled and overlapped display designs.

On the one hand, the depth multiplexing approach offers the best of both worlds: objects need not be tiled to be visible. Hence, ideally, less information is obscured. On the other hand, the potential for the content of one widget to interfere with another above or below it is introduced. This strategy raises questions about legibility, visual interference and what types of information this technique is best suited for. 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 acquire a more complete understanding of the constraints of such an approach.

Within the depth multiplexing strategy, we vary the level of transparency/opacity between the two interface items. *Fully opaque objects* reflect traditional window, palette, and menu design in current graphical user interfaces. *Fully transparent designs* reflect some of the unconventional interfaces such as those used in Heads Up Displays (HUDs) in aviation (Wickens et al., 1993; Larish and Wickens, 1991) or the displays used 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. S*emi-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). Semi-transparent interface examples are few and a structured evaluation methodology has not been applied previously.

Within our design space (Harrison et al, 1994), we wish to classify and evaluate a variety of semi-transparent interface widgets. Broadly defined these include menus (pull-down, pop-up, and radial or "pie" menus), palettes (tear off tool menus), dialogue boxes (especially interactive scrolling dialogues), windows, and help system screens. These widgets appear in many applications and at least temporarily obscure part of our work surface. The degree to which they persist (seconds versus minutes or hours) largely determines how disruptive they may be. In many situations, our primary task or the area of interest becomes the "background" layer while these objects appear in the foreground to enable us to carry out activities that are ultimately reflected in the now-obscured background layer. Transparent interfaces allow the user to observe these changes without obscuring our task layer. We do not need explicit window movement or manipulation strategies to switch attention between the menu or dialogue box and the underlying item of interest (which may be providing visual feedback based on our selections). For example, we may change font size, perform cut-and-paste operations, change colour of items, alter size of graphical drawing objects, or manipulate other attributes.

One additional advantage of transparent interfaces is the preservation of context or inherent task structure. For example, there are many examples of work scenarios where users work with two windows, two displays, or two versions of documents. One task is inherently active (or a foreground activity) while the other is passive (or a background activity). We may wish to preserve a sense of awareness of the background task while focusing on the foreground task. For instance, help screens are momentarily active while we read them and the context they are called from temporarily becomes a background task. In collaborative work, we often have a shared document (foreground) that we are co-authoring and a video connection or window to our colleague (for awareness – background). There are numerous such examples. Transparency allows us to preserve these relationships.

A number of design issues arise. How does transparency affect legibility and under which circumstances is it worst? Are there design techniques which will help us to visually separate the different layers and which information belongs to each layer? Which types of information improve or degrade this separation? We present our framework and methodologies for exploring these issues and answering some of these questions through a variety of evaluations.

THE CENTRAL ISSUE AND FRAMEWORK

The central issue in designing semi-transparent interfaces lies in supporting users' visual attention allocation. The very reason to make widgets semi-transparent is to enable the user to simultaneously monitor two different items or, observe one item while maintaining an awareness of the state of the other item or object. In other words, on the one hand we want to facilitate users dividing (or switching back and forth rapidly) their attention between two objects. On the other hand, sometimes we wish to observe a single object while ignoring all others. In other words, the design of the interface should also allow users to focus their attention on one object. In this cases, we need to minimize the degree to which the unimportant items interfere or accidentally attract our attention away from the desired object. Accommodating the conflicting requirements of focused and divided attention is often a critical issue in engineering psychology (Wickens, 1992). We are evaluating how partially transparent widgets facilitate divided attention while supporting focused attention and efficient attentional switching between objects.

Focused and Divided Attention

Our research investigates two types of attention: focused attention and divided attention. Focused attention is the ability to channel attention to one source of information in the environment. Problems in focused attention are the tendency we have to be distracted or unable to concentrate on this single source of information. Failures in focused attention result from the processing of unwanted information forced by environmental events (e.g., extraneous input such as light or noise or difficulties in legibility) despite the subject's best efforts to shut them out (Wickens, 1992). If there is too much visual interference from another object (particularly in the case of transparent overlapped objects) we may be unable to focus on either single object properly.

Divided attention, the other side of the coin, is the ability to attend tomultiple objects simultaneously. Failures in divided attention therefore are failures to divide our attentional efforts between the multiple sources of information (time-sharing), all of which we need to process. If the task requirements, stress levels, or workload are too demanding, we fail to divide our attentional resources adequately to process the multiple inputs. In transparent interfaces, we need to be able to clearly see both objects and to separate out the components that belong to each.

Divided and focused attention can be defined in terms of a multi-channel model (Treisman, 1969). Our attention can be focused on a channel of information where all events within this channel can be roughly processed in parallel. Simultaneous events within a channel must be processed in parallel; we cannot block one out. Information which appears in separate channels must be processed serially forcing us to block one channel while focusing on another.

A number of characteristics can be used to define a channel. In essence the more similar two "signals" or information sources seem to be, the more likely that they will be perceived as a single channel. For example, the pitch or tones of sounds, the gender of voices, the colour, size and font type of text, or the proximity of graphics – all define whether two items are perceived as part of a single information source. The most obvious property that defines a channel is space – proximal locations in the visual field. Hernandez-Peon (1964) first proposed the spotlight metaphor to distinguish the breadth of allocation of attention from the stability or detail of allocation. Attention was said to be "a beam of light in which the central brilliant part represents the focus, surrounded by a less intense fringe". The breadth of attention was defined by the width of the beam and the extent to which it "roamed" in scanning the visual field. Breadth typically includes some notion of the number and range of cues that fall within the beam (i.e., are integrated into a single precept). Although this theory has a number of flaws, it is still widely used to capture the basic concept of visual attention and peripheral awareness.

In summary, there are three critical attentional components: the ability to divide attention between two items, the ability to separate the visual characteristics of each source and focus on any single item with minimal interference from other items, and the switching cost (time, mechanism, learning, awareness) of shifting attention from one item to another.

Applying Attentional Theories to Transparent Interfaces

We wish to evaluate the objects and windows which appear at varying degrees of transparency. This will enable users to see-through the top-most interface objects to overlapped items which would normally be obscured in traditional graphical interfaces.

To facilitate focused attention (ignoring information from the background layer) we want to make the attributes of the information on foreground objects as different from the background as possible. There are many ways of achieving this (with varying success), such as different colours, content attributes – analogue (images or graphics) versus verbal (text based), font sizes or styles, etc. However, in most cases these dimensions are determined by the application tasks and therefore there is little freedom for us to manipulate these parameters. The degree of transparency between layers is one apparent dimension that can be changed to support focused attention. Low degrees of transparency (almost opaque) distinguishes the appearance of the foreground and background object, allowing the user easily focus attention on the foreground.

For divided attention (being able to see both foreground and background layers), a high degree of transparency is desirable to support higher visibility of both layers. However, the user must still be able to separate which features belong to the foreground and which to the background in order to accurately perceive the objects.

Clearly there is a trade-off between these two goals. We have characterized this trade-off in Figure 2 which provides a framework for this research. From this analysis, we can predict that the optimal degree of transparency is determined by the trade-off of supporting both focused and divided attention.

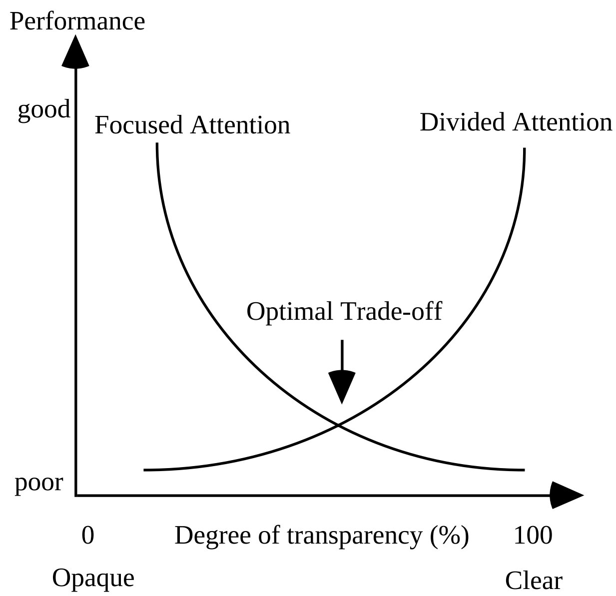

Figure 2. A simple model of transparency selection. As degree of transparency increases, it gets easier to divide attention between information on the top object and information on the background object but more difficult to focus attention on either object exclusively. The optimal transparency (OT) is a result of a trade-off. The curves and the location of optimal transparency in the figure are hypothetical but may reveal the trend. The non-linear nature of the curves is also proposed but appears to be supported from our preliminary experimental work.

RESEARCH METHODOLOGY AND PRELIMINARY RESULTS

We are taking *two* complementary approaches to study the silk widgets: formal experiments and realistic field studies.

To reveal how focused and divided attention exactly changes, i.e. how the curves in Figure 2 are shaped, we are conducting formal experimental studies with well controlled models and simulations. These experimental results provide us with precise measures on how well the user can see both foreground and background (particularly the background) information (i.e. divided attention) and on how high the interference is between the two "layers", by varying degree of semi-transparency in between the two layers.

However, we realize that controlled experimental paradigms address a restricted set of design dimensions only. The real applications consist of a much richer design space. Therefore, we have also developed several prototype systems which are more representative of real world applications. We are evaluating these systems and observing user behaviour to gain further insights into the design of transparent user interfaces. This combined research program allows us to further formulate research issues while remaining confident that our research results have external (real-world) validity. The two approaches are conducted in parallel and as iterative design evolutions.

The Formal Experiments

Our first set of formal experiments was used to evaluate the possible shape of the curves depicted in Figure 2 and to increase our understanding of how transparency interacts with divided and focused attention problems. To this end, we used a very simple but robust task to measure interference between two layers called the Stroop Effect (Stroop, 1935). There have been over 700 articles on various permutations of this effect – it is consistent, well documented, and cannot be voluntarily manipulated (MacLeod, 1991).

In traditional Stroop tasks, a series of words are presented in a chosen colour of ink (e.g., red, green, blue, and yellow ink). Subjects must name the *ink colour* while ignoring the word. Some of the words are neutral (e.g., uncle, shoe, cute, nail); other words are the names of conflicting colours (e.g., yellow, blue, green, red). Consistent, significant performance degradation occurs when conflicting colour words are used and subjects attempt to name the colour of the ink (e.g., the word "red" appears in green ink; the correct response is green). (See Figure 3, colour plate.)

For our experiment, we used a coloured rectangle on the foreground layer and the word appeared through the colour patch on the background layer. We varied the transparency of the colour patch such that the word appeared at varying levels of clarity. At high levels of transparency (e.g., 100% - clear) we anticipate that users will experience high levels of interference from the word when they try to name the colour (difficulty in focused attention). As the colour patch becomes more opaque the interference from the word should decrease (making focused attention easier). This would support the focused attention curve proposed in Figure 2.

To test the divided attention curve proposed we used the second component of the Stroop Test: word naming. In this case users are asked to ignore the colour patch and read the word in the background layer. This experiment reflects more of a legibility test, necessary for divided attention. The colour patch in the foreground is always clearly visible and perceived. By reading the background word the user is, in effect, creating a divided attention task. Since we are interested in the design trade-offs for transparency versus awareness of the background layer, this experimental task reasonably approximates our goal, though it is an extreme case. (By using an extreme case we should be able to test for threshold values which are more constraining that those of the real-world tasks.) At high levels of transparency (e.g., 100% - clear) it should be very easy to read the background word (divided attention is easy). At more lower levels (opacity increases) it should become progressively more difficult or impossible to read the word (loss of ability to divide attention).

For our experiments we used 4 colours: red, blue, green, and yellow. Words (helvetica, 78 point, uppercase) appeared "through" the coloured rectangular patch. We used neutral words UNCLE, NAIL, CUTE, and FOOD in addition to the four colour names. Transparency levels were varied as: 0% (opaque), 5%, 10%, 20%, 50%, 100% (clear). Images were produced using alpha blending on an SGI Indigo™ to ensure accurate transparency levels.

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 colour (a conflicting colour word was present), congruent colour (the colour word matched the colour of the patch), and baseline (colour naming only with no word present or word naming only with no colour present). 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. A total of 16 students from the University of Toronto were run as subjects[1].

Results

We found that decreases in transparency level (i.e., more opaque) did indeed dilute the Stroop Effect in a statistically significant way ($F = 1.3$, $p < .05$) for our focused attention task. In this sense performance improved and error rates were lower. This seems to support the proposed focused attention curve depicted in Figure 2. A post-hoc comparison of means shows that levels of 5% and 10% transparency are statistically indistinguishable from 0%, suggesting a minimum cut off point in performance (which we also hypothesized). As the levels increase beyond 20% the interference effects become more pronounced. Faster performance times were considered better ratings per Figure 2. The mean response times were: 0% - 603 ms, 5% - 624 ms, 10% - 624 ms, 20% - 680 ms, 50% - 687 ms, 100% - 713 ms.

For the word naming task (divided attention) results also seem to support the proposed curve in Figure 2. Word naming seems highly error prone at levels of 5%. At 10%, about 1/3 of the trials were not properly visible to the subjects. (Also at the 10% level, subjects seemed to perform slightly better, depending upon what the background colour was.) At 20% levels and higher, all words were easily read though the response times improved slightly as transparency increased. Transparency significantly effected both errors and performance ($F=2.37$, $p< .001$). These results suggest a cut-off point between 5% and 10%, where an unacceptable number of errors and poor performance occur. Our results also suggest that for transparency levels between 50% and 100% performance is very high and improvements with increases in transparency are marginal.

For lower levels of transparency where the word was only slightly visible, yellow and green hues seemed "easier" to detect the word on (as reflected in the post-hoc analysis of means). 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 colours

Our experiments suggest promising results which relate our theoretical framework to empirical data. However, they were conducted under simplistic, controlled task scenarios. Clearly we must conduct further research to determine how well our theoretical framework fits more realistic and varied task scenarios. To assess this we have built more representative prototypes for evaluation purposes.

The Informal Studies - Design and Observation Through Prototyping

Together with several other prototypes, we have developed a test-bed program, SilkWidgets, on SGI™ workstations (Figure 1, colour plate). The prototyping process provides us with the opportunity to observe the behaviour of silk widgets in various practical scenarios, to explore the design space of silk widgets, such as the degree of transparency and the design of fonts for silk widgets, and to further hypothesize research issues for formal experimentation.

The Prototype SilkWidgets

SilkWidgets is a colour drawing program which contains elemental functions that are often used in end-user software packages. Text, lines, circles, blocks and

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

other simple geometric forms can be drawn either opaquely or semi-transparently. Image files can be imported or exported. Pull-down and pop-up menus are incorporated and the level of transparency of these can be adjusted. We also have a text-based help screen which can be set to varying levels of transparency.

Informal Observations

Many of our observations from prototype usage are in agreement with our theoretical analysis. For instance, when two sources of information are distant in coding format, e.g. the foreground is textual (menu) and the background is analogical (graphics), the user had little difficulty in focusing attention on the foreground information even though the widgets are highly transparent and the background is clearly visible (e.g., Figure 4a, colour plate). In contrast, when both the widgets and the background have textual information, it is more difficult to focus attention on one source and ignore the interference from the other. Differences in font, colour and size improves the separation of the two sources (Figure 4b, colour plate). Preliminary testing also indicates that increases in expertise improves performance. In particular, our expert menu users select items based on spatial location and did not require high legibility. Also users who were highly familiar with the background content seem to have less difficulty separating information into the correct layers.

We have tested different degrees of transparency with a variety of backgrounds. We found that the optimal degree of transparency depends on the information formats of the background and foreground as discussed earlier. When the background information format is determined by the application task and the designer/user has no choice on it, 20%-40% of transparency appears to work well on most backgrounds. Within that range, both background and foreground objects are visible (divided attention supported) and yet the contrast difference in background and foreground objects is large enough to enable the user to separate out the two widgets and focus on one (usually the foreground) when needed (focused attention supported). This again is in approximate agreement with our formal experimental studies. (The Stroop test identifies the highest constraint which is the lower bound – 20%.) Figure 5 (colour plate) shows a few examples of the silk widgets with different levels of transparency.

Visibility and Focused Attention

One issue closely related to supporting focused attention is the visibility of the text on the silk widgets. When text-based items such as menu choices appear superimposed over coloured backgrounds, they may become difficult to read and therefore it is impossible to focus attention on the foreground information. This is particularly problematic when the background colour matches (or approaches) the font colour (Figure 6a, colour plate). Since any colour may appear in the background (created by the application user), this can happen no matter which colour is chosen for the text on the widgets (by the application designer).

To address this, we have developed "anti-interference" (AI) fonts (Figure 6b, colour plate). It is well known that contrast is the most critical dimension of visibility. In AI fonts, the opposing outlines of the text are rendered in a colour which has the maximal contrast to the colour of the text. For any selected text colour vector [R, G, B], our AI font algorithm calculates the luminance value Y according to the YIQ colour model used in television broadcasting (Foley et al 1990, page 589). Note that the red, green and blue components are not equally weighted in contributing to luminance.

$$\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.275 & -0.321 \\ 0.212 & -0.528 & 0.311 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

Based on the value of Y, our algorithm then determines the outline colour with the maximal luminance contrast. In practice, only two colour vectors can be the candidates for the solution: [0,0,0] (black) when $Y > Y_{max}/2$ or $[R_{max}, G_{max}, B_{max}]$ when $Y < Y_{max}/2$, where Y_{max} is the maximum luminance value and $R_{max}, G_{max}, B_{max}$ are the maximum red, green and blue value respectively.

Since an AI font has two opposing colour components, it remains visible in any colour background. Figure 6b shows the worst case scenarios when the background colour is the same as the text or the text outline. As we can see, the AI font is "impenetrable" by any background colour.

Future Design Variations

We are currently exploring design variations of silk widgets based on our experience from both the formal experiments and the prototype usage observations. We are also learning from the designs of television captioning. One such design variation is to apply varied degrees of transparency to different parts of a widget (within a single widget) according to the information density. Figure 7 (colour plate) illustrates this idea to some extent. Lower degrees of transparency (high opacity) are chosen for parts of the widget where the information is dense and more transparency is used on parts of the widget where information is sparse. This may facilitate perception of information in the background. Interestingly, our visual systems tend to group the objects on the widget even when the transparency is uneven so that focused attention may still be well supported. 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).

CONCLUSIONS AND FUTURE DIRECTIONS

We have proposed a framework for investigating transparent user interfaces, based on theories from divided and focused attention. Our results thus far show promising support for our hypotheses about the interaction between these two forms of attention and also for determining a potential trade-off point or optimal design. We are continuing with our formal experimentation to further relate our theoretical framework more closely to possible design dimensions.

We are also continuing evaluations using our "silk widgets" prototype to gain insights into possible design permutations – particularly those which are application independent and can be globally implemented across object types. Thus far several promising design solutions have emerged. We wish to evaluate a number of other user interface mechanisms which seem to facilitate the use and application of transparency. This will help us to ultimately derive design guidelines. We are also attempting to categorize critical information characteristics which will influence the success of transparent interface implementation. For example, superimposing text over text, text over images, text over line drawings, etc. These design ideas are being incorporated into progressively more realistic application domains.

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. Our results thus far show promising advantages for creating new user interfaces and interaction techniques using transparency. We are exploiting possibilities of new technology in a way that is sensitive to both psychological and task constraints.

ACKNOWLEDGEMENTS

Support for this research is gratefully acknowledged from the Natural Sciences and Engineering Research Council (NSERC) and Alias Research Inc. This work has been significantly improved with the support from many of our colleagues. We would particularly like to thank Dr. Hiroshi Ishii of NTT, Dr. Gordon Kurtenbach and Steve Spenceley of Alias Research Inc, Dr. Chris Wickens of University of Illinois at Urbana-Champaign, Anu Rastogi, Dr. Paul Milgram and Dr. Colin MacLeod of University of Toronto for their insights and technical assistance.

NOTE

The colour plates have been rescaled and printed on a colour printer. For final editions, we will use high quality slides to produce the final colour plates. This will produce smaller images with better clarity. The actual final length of this paper will therefore be shorter.

REFERENCES

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.

Foley, J.D., van Dam, A., Feiner, S.K., and Hughes, J.F., (1992), *Computer Graphics - Principles and Practice* (second edition), Reading, Massachusetts: Addison-Wesley.

Hernandez-Peon (1964). Psychiatric implications of neurophysiological research. *Bulletin of the Menninger Clinic*, 28, pp. 165-185.

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.

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.

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).

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.

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

Treisman, A. M. (1969). Strategies and models of selective attention. *Psychological Review*, 76, p. 282-299.

Wickens, C. D. (1992). *Engineering Psychology and Human Performance*, (Second Edition) Harper Collines Publishers.

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.

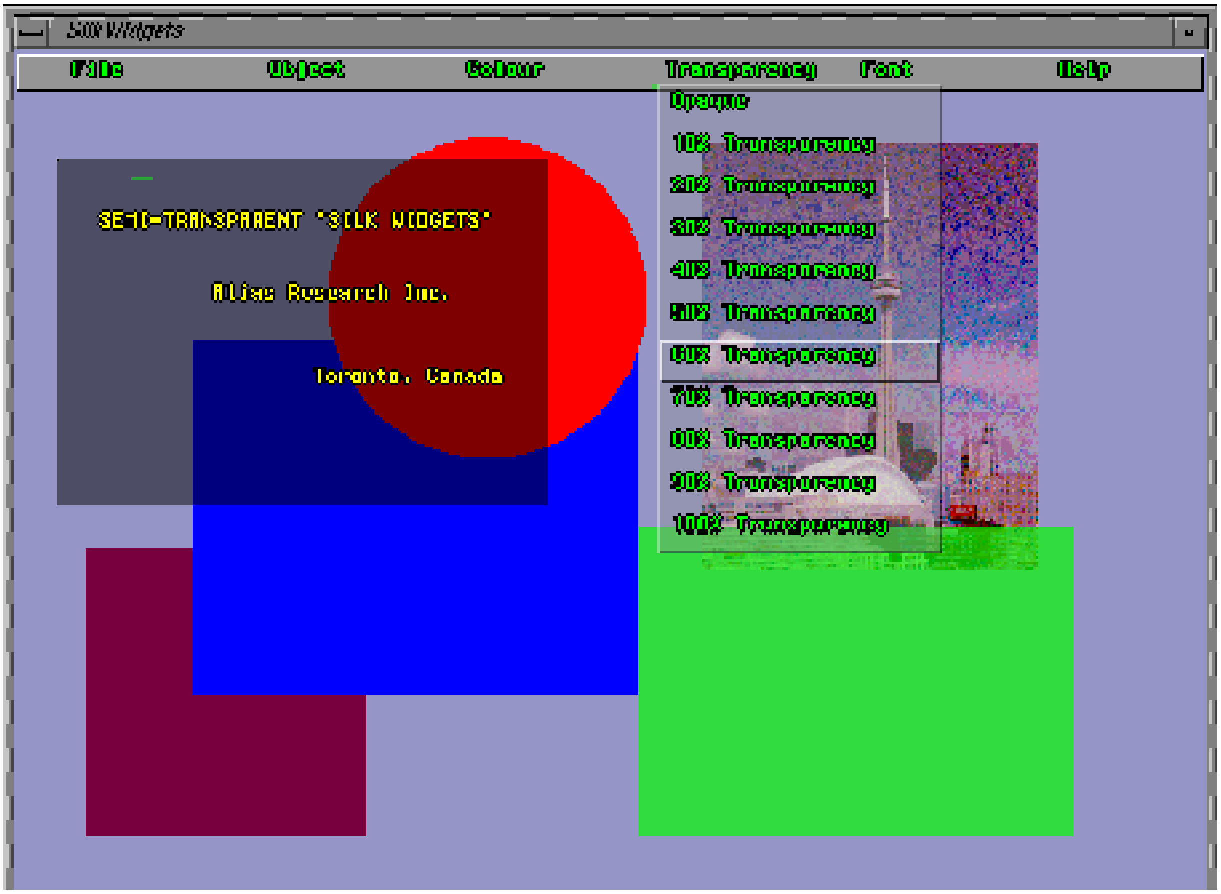

Figure 1. SilkWidgets - the pull-down menu is semi-transparent

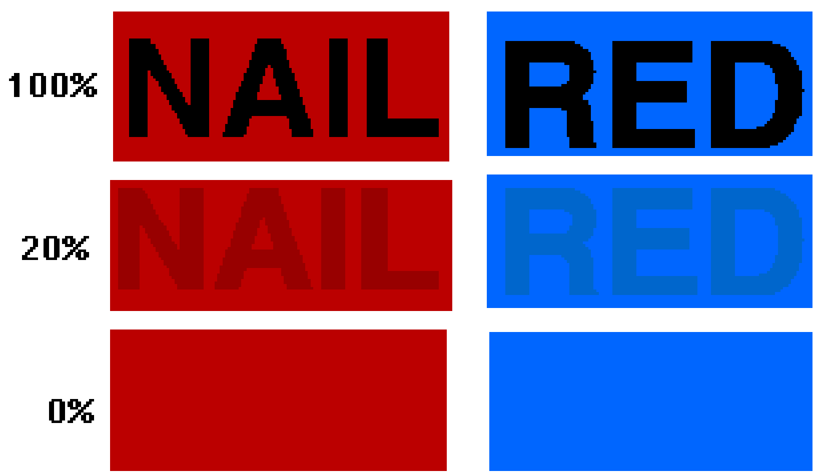

Figure 3. Sample Stroop Test data. The column on the left contains neutral words, while the rightmost column shows color conflict words. The two tasks are to name the color patch – ignoring the word, and name the word – ignoring the color patch. A subset of the transparency levels used are shown here as an example.

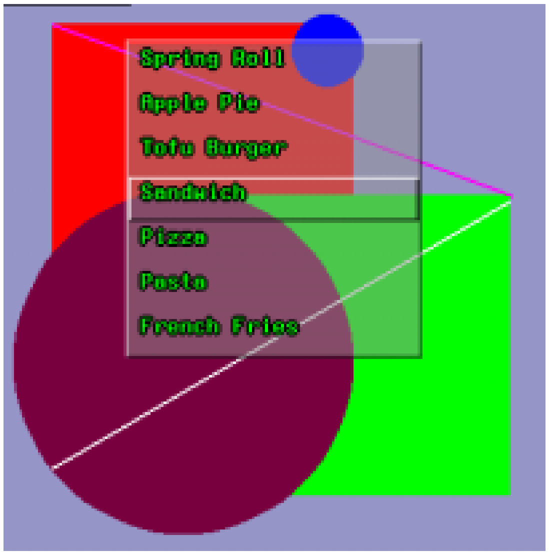

Figure 4 (a). Pop-up silk menu on graphics

Figure 4 (b). pop-up silk menu on text

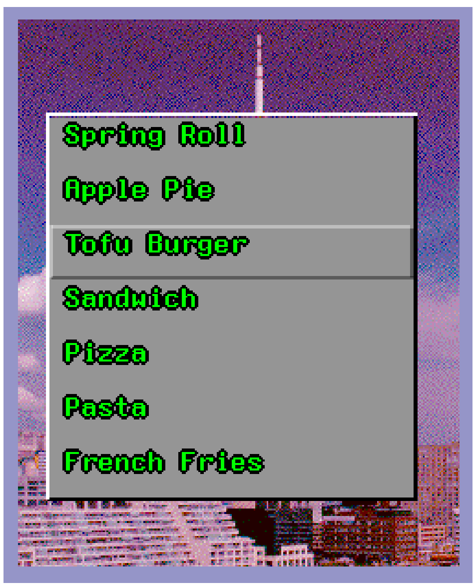

Figure 5 (a). Pop-up silk menu at 0%

Figure 5 (b). 20% transparent

Figure 5 (c). 40% transparent

Figure 5 (d). 80% transparent

Figure 5 (e). 100% transparent

Figure 6 (a). Plain font style

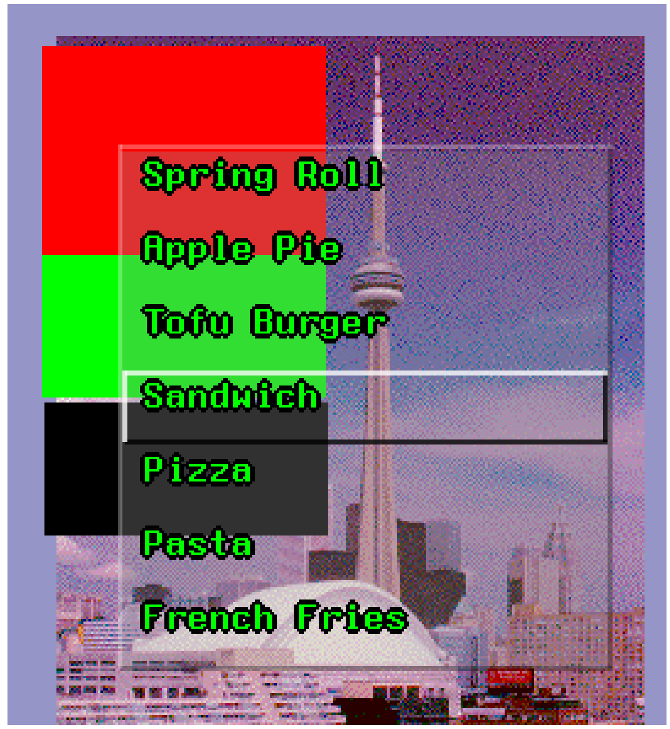

Figure 6 (b). "Anti-interference" (AI) font style

Figure 7, "Silk" Widget of varying transparency on television sportscast