Readability of a background map layer under a semi-transparent foreground layer

Saturnino Luz
School of Computer Science and Statistics
Trinity College Dublin
Dublin, Ireland
luzs@cs.tcd.ie

Masood Masoodian
Department of Computer Science
The University of Waikato
Hamilton, New Zealand
masood@cs.waikato.ac.nz

ABSTRACT

This study investigates the readability (interpretability) of information presented on a geographical map onto which a semi-transparent multivariate selection layer has been overlaid. The investigation is based on an information visualization prototype developed for a mobile platform (tablet devices) which aimed at supporting epidemiologists and medical staff in field data collection and epidemiological interpretation tasks. Different factors are analysed under varying transparency (alpha blending) levels, including: map interpretation task (covering "seeing map" and "reading map" tasks), legend symbol and map area type. Our results complement other studies that focused on the readability characteristics of items displayed on semi-transparent foreground layers developed in the context of "toolglass" interfaces. The implications of these results to the usability of transparency variable selection layers in geographical map applications are also discussed.

Categories and Subject Descriptors

H.5.m. [**Information Interfaces and Presentation (e.g. HCI)**]: Miscellaneous

General Terms

Human Factors

Keywords

Information Visualization, Transparency, Geographical Maps, Epidemiology

1. INTRODUCTION

Maps, in one form or another, have existed since the dawn of human civilization, with the more widespread use of geographical maps beginning with the age of exploration around the 15th century. The first use of maps for visualization

of geospatial data (disease maps) dates back to the 18th century [11]. In the 21st century, maps have gone beyond their traditional role of presenting data and can now be seen as flexible interfaces to geospatial data [12]. In this context of cartographic information visualization, according to MacEachren and Kraak, maps need to support "information exploration and knowledge construction", where "interaction" with "dynamic maps" allows the user to explore the map "without hypotheses about the data, and [where] visualization tools assist in an interactive, unencumbered search for structures and trends, with one goal being to prompt hypotheses" [16].

As the use of maps to explore geospatial data is not new, interactive system design can benefit (and have benefited) from traditional practices. Bertin [2], in his *Semiology of Graphics*, for instance, uses the example of multiple maps to show that by juxtaposing separate images on the 2D plane the viewer can utilize their ability to attend to a particular location, which is easier than attending to a particular shape in different locations, to see patterns in geospatial data more readily. Similarly Tufte [19] presents the idea of "small multiples", as a means of allowing exploration of quantitative data on the 2D "flatland" of printed paper and computer displays. More recently micromaps [3] have been proposed as a useful technique for highlighting "geographical patterns" and "association among the variables" in data.

It should be noted that what Bertin and Tufte propose is really an extension of the side-by-side map comparison technique, as a way of investigating patterns between maps. This method is hindered by the fact that human vision is not generally very effective in judging spatial correspondence between patterns of variables on side-by-side maps [15]. To deal with this problem, Bertin [2] proposes the use of a matrix arrangement of maps images in which the row-column position of the maps can be changed to allow better grouping of maps according to their dominant spatial patterns. However, Bertin's technique was devised in a time that predated the use of interactive computer. Since then, more interactive solutions that build on these basic ideas have been proposed for computer-based visualizations. A good example of such developments is the use of interactive micromaps, introduced by by Carr and Pickle [3].

The development of web-based map systems over the past decade has made it possible to utilize the API provided by such systems to overlay geospatial data on these interactive maps. This is generally achieved by using one or more *layers* of data, which can be individually turned on and off to show

or hide pre-defined sets of data. Most of these systems are however targeted towards users who need to access specific information (e.g. all the restaurants in the selected area). Furthermore, they are also mainly useful for cases where the user has a small number of layers for maps and data with which they are familiar [5].

Our focus, on the other hand, has been on providing interactive map visualization techniques for exploration of geospatial data, to allow the user to discover patterns and trends that are not known in advance. This is particularly relevant to the area of disease surveillance, where geographical patterns of disease spread need to be explored by manipulating and selecting a large range of recorded disease data attributes. To make this possible, we have developed a prototype system for mobile devices (tablets in particular) which is intended to be used in the field by healthcare professionals and epidemiologists for data collection, and to support the work of epidemiologists in monitoring the spread of diseases [14].

The large number of attributes and variables that are characteristic of epidemiology applications, and the need to allow the user to manipulate the ranges and values of such variables and attributes while observing changes on the map pose challenges for the design of suitable interfaces, specially where the available screen space is limited. The use of transparency layers has been explored as a means of combining an attribute setting mechanism with visualization of a background working area, both as transparent tools to alter the properties of specific background objects [7] and as "lenses" for map exploration [13, 5].

In this paper we present a user study conducted to evaluate the readability, or what could be more accurately defined as the *interpretability*, of geographical and categorical information about cases of disease. This information is presented on the background layer of a map, with user interface controls are presented on a semi-transparent foreground layer. Although the application area of our study is more specifically related to geospatial disease maps, and focuses on readability of information on the background layer, it complements other existing studies which have examined the readability of information on the semi-transparent foreground layers.

2. RELATED WORK

While there is a sizeable literature on human perception of transparency, the research reviewed in this section covers studies of human performance on user interfaces that employ transparent or translucent elements. As our focus here is on investigations on the use of transparency in user interface elements (in the context of human computer interaction), we have not included, for instance, studies that examine the use of transparency in medical imaging. While informative, the latter tend to focus on volume perception issues, such as depth perception [9, 10], rather than two dimensional data interpretation tasks.

Harrison et al. [8] tested the legibility of semi-transparent icons and text on the foreground of user interface screens, but not the legibility (or intelligibility, given that users usually do more than read) of foreground elements given a semi-transparent foreground. They identified three attention factors that might affect the users' interaction with transparent interfaces, namely:

(1) one's ability to divide one's attention between two item classes, such as between foreground and background objects;

(2) one's ability to separate different sources, such as background and foreground, with minimal interference (by "interference" we mean such things as possible confusion about which items belong to the foreground and which belong to the background);

(3) and the cost of switching focus from one item to another.

We propose that item (2) could be further specified by distinguishing between:

(2a) user's focus on the foreground, and

(2b) user's focus on the background.

From this perspective, the work reported by Harrison et al. [8] only examines (2a). They found that icon type (text, solid image or wireframe), background type (ditto) and transparency level all significantly affect user performance. In most cases, however, a 50% opacity (α blending level) works nearly as well as 100% opacity (i.e. no transparency). Performance decreases substantially (that is, the response time doubles) with $\alpha = 50\%$ and deteriorates further with $\alpha = 10\%$. Wireframe backgrounds give worse performance than text backgrounds as transparency increases. From the perspective of the map-based study that we will present here, this means that locating a place on the map by reading its name should be easier than locating delineated regions, if those regions form complex shapes.

Similarly Cox et al. [4] have found that people work best with a 50-75% transparency level on tasks that involve attention distribution between background and foreground. The difference with respect to [8] is that Cox et al. have tested full overlays rather than just "toolglasses".

In a related work, Harrison et al.[7] have used a Stroop effect test, which consists in asking the subject to name a colour or a word, sometimes adding linguistic interference, such as showing the word "green" on a background seen through a red tinted foreground. In this work they aimed to assess item (2a) but also, in part, item (2b). The interference of foreground on background (2b) was assessed by asking the subjects to read a word on the background through a coloured foreground layer (4 different colours were tested: blue, yellow, green and red) of varied opacity levels. Interference of background on foreground (2a) was assessed by asking the user to name the color of the foreground. As transparency levels increased, performance (response time) degraded on 2a and improved on 2b, as expected. Although their evaluation of 2b was limited to word reading, they did find significant differences in performance between 10% and 20% transparency, and a marked increase in performance for 50% transparency and above.

Gutwin et al. [6] also tested foreground visibility (2a). They assessed the effects of transparency on the user's ability to select "floating" palettes and menu windows over a dynamic background when the transparency level of the foreground object is variable (near transparent when the cursor is away; near opaque when the cursor is over it). They found that (1) targeting cost increases with transparency but the increase is rather minor for values up to 75%, and that (2)

complex background images require lower transparency values in order to preserve targeting performance. They tested the following variables (3x4x2x2 design): *selection* (3 levels: one-d, multi-d, and palette object selection), *transparency level* (4 levels: 100%, 90%, 75% and 0%), *background complexity* (simple vs. complex) and *transition style* (sharp vs. gradual). The study showed interactions among all levels of these variables, but the effect sizes seem quite small, specially between intermediary transparency levels (75% vs. 90%). For instance, the maximum difference between the 0% and the 100% transparency conditions for the multi-d task was only 23% (i.e. with transparency set to 100% users were 23% slower than when targeting fully opaque objects). For the palette task, the maximum difference was only 8%. The main finding was that users stated that the maximum usable transparency for complex backgrounds is 75% (90% for blank backgrounds).

As regards transparent layers for map visualization on mobile devices used in navigation tasks, McGookin et al. [17] experiment with a transparent layer containing a photo of a "point of interest" (POI) overlaid on a map. The goal is to help the user locate that point of interest. In their interface, the transparency can be varied either manually or automatically according to the distance to the POI. The assumption is that as the user approaches the POI they will be more interested in seeing the image than in seeing the map, and therefore transparency should decrease. Three conditions were tested: (1) no transparency layer (control), where the image was accessed as a separate screen by tapping, (2) manual transparency, where the image is shown overlaid on the map with its transparency level adjusted manually by the user, and (3) automatic transparency, where transparency decreases as the user approaches the POI. Although the results suggested that transparency may lead to better task performance in locating the POI, no statistical significance was found.

Two other works that investigate the use of transparency in maps are by Lieberman [13] and Elias et al. [5]. The latter is interesting because it proposes a taxonomy of tasks for map users (the word "users" here refer to the people they interviewed, namely, experts such as map librarians, GIS staff and students). In their taxonomy people correlate maps by *familiarization* (i.e. relating features of a map to contexts that are well understood by the user), *evolution* (i.e. detecting changes over time) and *fusion* (synthesis of two maps from categories unfamiliar to the user). They derive some design principles from these tasks and present an implementation that illustrates some of these principles. However, they present no evaluation of their implementation.

Lieberman [13], on the other hand, superimposed a zoomed view of part of the map on its overview (i.e. an overview plus detail technique). This is somewhat closer to the user interface concept we have studied, since the transparent overlay covers the entire map and the users are (at least some of the time) interested in interpreting the information contained on the background layer. Once again, as with the the work by Elias et al., Lieberman [13] does not present a formal evaluation of the proposed technique.

3. USER STUDY

The studies reviewed above have broadly established that users can manipulate foreground layer tools reasonably effectively at transparency levels of up to 75%. In complement

to this, we decided to assess the effect of different levels of foreground transparency on the user's ability to interpret content displayed on the *background* layer, while performing geographical map related activities on the background layer.

Visual tasks commonly associated with map interpretation are reading (e.g. names of towns, rivers, roads etc) and legend and symbol interpretation (e.g. location, grouping and identifying overall patterns of sets of symbols on the map). According to Bertin [2], there are two types of maps: "seeing maps" and "reading maps". Seeing maps allow the map user to answer questions about the overall data set (often at a glance) as represented by the legend symbols. These include questions of attribute grouping and geographical location. Reading maps, on the other hand, require the map user to examine all symbols at the elementary level in order to find attribute relations. The latter usually comprise comprehensive superimposition of symbols.

Another source of variation in map interpretation tasks involve the placement of symbols on different regions, represented on the map by different levels of visual complexity. Examples include urban areas, which contain complex information and are thus visually "busy", and rural areas, which are usually represented by uniform and continuously shaded shapes.

Our experimental design aims to cover reading versus seeing types of tasks on both urban and rural areas of interest, over varying foreground transparency levels. We expect user performance to degrade, and perceived difficulty in performing the task to increase, for background reading tasks as transparency decreases. We thus seek to reject the null hypothesis

H1: response time and difficulty ratings remain the same for different levels of foreground transparency in background layer interpretation tasks

Conversely, we expect performance to improve, and perceived difficulty to decrease, as transparency decreases for foreground layer interpretation tasks. The null hypothesis in this case, call it **H2**, is analogous to H1. We also expect reading maps to be considerably harder to interpret than seeing maps, and that decreasing transparency should increase the difficulty of the former with respect to the latter. The null hypothesis to be rejected in this case can be stated as

H3: response time and difficulty ratings are the same for reading maps and seeing maps, for all transparency levels.

Lastly, we anticipate that the more complex backgrounds of urban areas will make both the background and the foreground interpretation tasks harder to accomplish, and that this difficulty will be accentuated by decreasing transparency (background task) or increasing transparency (foreground case). Therefore we have the following null hypotheses to reject:

H4a: in the background reading task, response time and difficulty ratings are the same (for all transparency levels) regardless of whether the symbols are placed on urban or rural areas.

H4b: in the foreground reading task, response time and difficulty ratings remain the same (for all transparency levels) regardless of whether the elements on the foreground layer are placed over urban or rural areas.
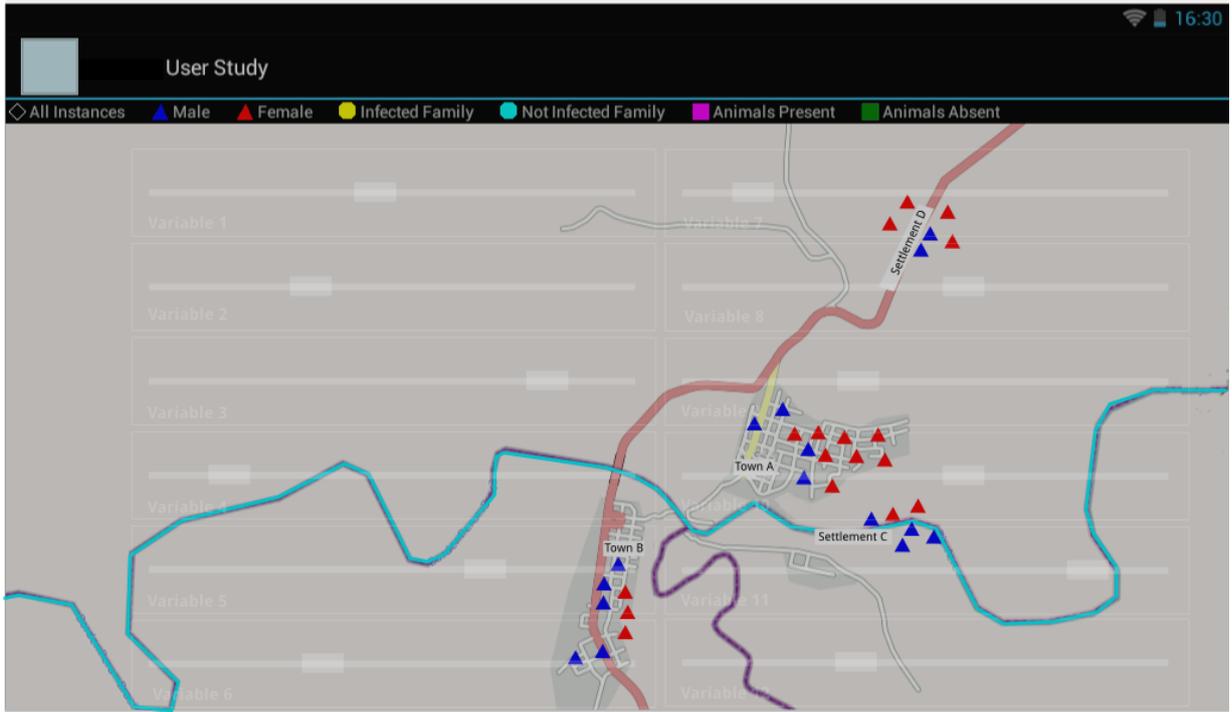
Figure 1: Sample image used in the experiment, displaying a 75% transparent foreground layer over a map background. In this case, symbols indicate the gender of each patient case.

3.1 Method

We employed a $3 \times 4 \times 2$ experimental design that comprises three factors, ranging over the following possible levels:

- three transparency levels ($\alpha \in \{25\%, 50\%, 75\%\}$) for the foreground layer,

- four binary legend symbol classes (gender, family infection, presence of animals, 3 symbols combined) for presenting individual disease case data,

- four map areas (small town, large town, river side, road side), which we grouped into two overall types: urban and rural areas

Transparency level and legend symbol were treated as within-subject factors (i.e. each subject experienced all levels of each factor) while map areas were treated as between-subject factors. Legend symbols placed on a map could be shown individually (seeing maps), or as superimposed symbols (reading maps). The different classes and values of these symbols are shown in Table 1. These symbols, from left to right, represent the gender of the patient, whether any of their family members also have the disease, whether they have animals in their houses, and various combination of these values.

Table 1: Legend symbol classes

Gender		Fam. Inf.		Animal		Combined
male	female	yes	no	yes	no	(example)
▲	▲	●	●	■	■	▲

The experiment started with a brief tutorial presented to the study participants. The tutorial explained the study task using an example image consisting of a background map layer and a semi-transparent foreground user interface layer. The tutorial also described different types of map legends used in the experiment to represent individual attributes, as well as combinations of attributes of disease case data. The disease data set used for the tutorial was different to the one used for the actual experiment. After the completion of the tutorial the participants were directed to the trial.

For each study task question the participants were presented with an image consisting of two layers: a foreground user interface layer (with varying levels of transparency) showing a number of sliders (e.g. Variable 1, Variable 2 etc), and a background map layer on which various legend showing individual disease cases were placed across four different location groups (two urban areas: Town A and Town B, and two rural area: Settlement C, and Settlement D). Figure 1 shows an example image used in the experiment. In this image, for instance, two male and three female patient cases are shown in Settlement D.

The study task questions were divided into five groups, with four of them relating to information presented on the background layer map, and one relating to the user interface sliders presented on the foreground layer. Table 2 shows the questions used in our study.

In each of the four groups of questions relating to the background (groups 1-4), each question was presented to different participants with one of the 3 different foreground transparency levels ($\alpha \in 25\%, 50\%, 75\%$) and a fourth image without a foreground layer (i.e. only the background map). All the image/question combinations were presented randomly, and were counter-balanced. Questions in groups

Table 2: Task questions used in the study.

Group	Questions
1	In Town A, the majority of cases occur in households where ANIMALS are PRESENT.
	In Town B, the majority of cases occur in households where ANIMALS are PRESENT.
	In Settlement C, the majority of cases occur in households where ANIMALS are PRESENT.
	In Settlement D, the majority of cases occur in households where ANIMALS are PRESENT.
2	In Town A, the majority of cases occur in households where the FAMILY is INFECTED.
	In Town B, the majority of cases occur in households where the FAMILY is INFECTED.
	In Settlement C, the majority of cases occur in households where the FAMILY is INFECTED.
	In Settlement D, the majority of cases occur in households where the FAMILY is INFECTED.
3	In Town A, the majority of cases are FEMALE patients.
	In Town B, the majority of cases are FEMALE patients.
	In Settlement C, the majority of cases are FEMALE patients.
	In Settlement D, the majority of cases are FEMALE patients.
4	In Town A, all MALE patients are in NON-INFECTED FAMILIES with NO ANIMALS.
	In Town B, all FEMALE patients are in INFECTED FAMILIES with ANIMALS PRESENT.
	In Settlement C, all FEMALE patients are in INFECTED FAMILIES with NO ANIMALS.
	In Settlement D, all MALE patients are in INFECTED FAMILIES with ANIMALS PRESENT.
5	The slider for Variable 10 is set higher than the slider for Variable 7.
	The slider for Variable 8 is set lower than the slider for Variable 12.
	The slider for Variable 9 is set lower than the slider for Variable 11.
	The slider for Variable 1 is set lower than the slider for Variable 4.
	The slider for Variable 2 is set higher than the slider for Variable 5.
	The slider for Variable 3 is set higher than the slider for Variable 6.

1-3 are categorized as seeing map tasks, while those in group 4 are reading map tasks.

Questions in group 5 are, on the other hand, related to the foreground layer. These questions were not used to replicate what has been investigated in previous studies, but rather, to bring to the attention of the study participants the need for the foreground layer in performing map related tasks. As such, we did not use an extensive set of foreground task questions. In this group, questions 1 and 4 were used with foreground transparency level of 25%, questions 2 and 5 with 50%, and questions 3 and 6 with 75%.

All questions had yes-or-no answers. Once the participant chose a response to a study question, the system recorded the time taken to answer the question, and then presented the participant with a request to rate the difficulty of the question they had just answered on a Likert scale ranging from 1 (easy) to 7 (difficult). The participants were asked to respond to each question as quickly as possible, and were informed that the time taken to rate the difficulty level of questions would not count towards their task completion time.

Additionally, the trial included an extra question where the participants were asked to choose the combination of background and foreground that they preferred overall. The choices were foreground transparency levels 25%, 50% and 75% with the same background map layer.

A custom-designed interface was implemented for the experiment in Java. The system ensured that the order of presentation of the questions was properly randomized, that the between-subject questionnaires were distributed as evenly as possible among the study participants, and that the answers and response times (time to answer each question) were precisely logged. All images were all displayed at a fixed resolution of 1026×600 pixels. At the end of the experiment, the log files were processed by an R script[1] and collated into a single data set for statistical analysis (also done in R).

3.2 Participants

Once the experiment received approval from the Ethics Committee of Trinity College Dublin, a call for participation was circulated via email. A total of 22 people participated in this study. Of these participants, 10 were students, 8 were academics and 4 had other occupations. There were 17 male and 5 female participants. The age distributions were as follows: 9 were between 20 and 29 years of age, 6 were between 30 and 39, 5 between 40 and 49 and 2 between 50 and 60. None of the participants was colour blind, and 7 of them wore glasses.

3.3 Results

As with other experiments of this kind (e.g. [8]) accuracy was very high across all questions and participants (approximately 95% overall). Therefore we followed the commonly adopted practice of focusing on the response time as the main objective performance measure, complemented by the subjective difficulty ratings for each study task question.

The performance results were analysed through two-way repeated measures analysis of variance (ANOVA). The dependent variable (time) was log transformed and the result was found to be approximately normally distributed. The

---

[1]http://cran.r-project.org/
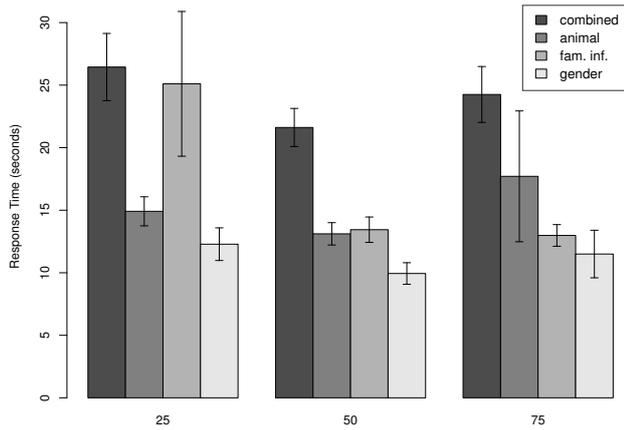
Figure 2: Response times for background map layer interpretation tasks, by legend symbol type for different foreground transparency levels (25%, 50% and 75%).
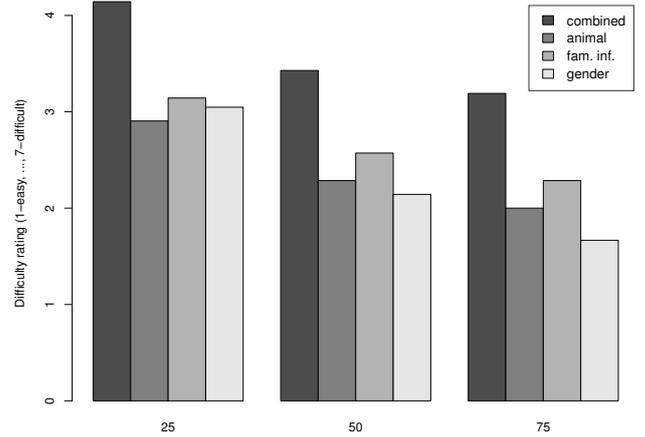

Figure 3: Difficulty ratings (Likert scale, 1 to 7) for background map layer interpretation tasks, by legend symbol type for different foreground transparency levels (25%, 50% and 75%).

Levene test did not point to any violation of the homogeneity of variance assumption ($F[11, 240] = 1.1, p = 0.39$, for the background task, and $F[2, 123] = 2.3, p = 0.1$ for the foreground task). Questions about background (designed to detect the effect of foreground transparency on background interpretability) were analysed separately from the questions about foreground items.

3.3.1 Background interpretation task

For this task, we found statistically significant main effects for *transparency* ($F[2, 240] = 4.9, p < 0.01$, and effect size $\eta^2 = 0.04$) and *legend symbol* ($F[3, 240] = 33.7, p < 0.01, \eta^2 = 0.28$). No significant interaction of transparency and legend symbol was found ($F[6, 240] = 0.69, p = 0.65$). Pairwise comparisons of the means using Tukey's Honestly Significant Difference (HSD) procedure indicated significant differences between transparency levels 25% and 50% ($p < 0.05$), between each of the individual legend symbols and the composite symbols ($p < 0.01$) and, somewhat surprisingly, between the symbol for gender, represented as a triangle, and the other two individual symbols, represented as a rectangle or a circle ($p < 0.01$).

The mean response times for each transparency level for this task, grouped by legend symbol type, are shown in Figure 2. In aggregate, as expected, the reading map tasks (represented by the dark grey, leftmost bars of each grouping, labelled "combined" in Figure 2) were also found to take significantly more time than the seeing map ones (Tukey HSD, $p < 0.01$) but no significant interaction with transparency level was found ($p = 0.75$). Finally, the region (rural vs. urban) on which the symbols were placed had a significant effect on performance, with questions based on rural backgrounds taking on average 15.9s to answer, while the ones based on urban backgrounds took 18.6s on average ($F[1, 19] = 4.16, p < 0.05, \eta^2 = 0.05$).

The subjective difficulty ratings of the background tasks provided by the participants also exhibited significant effects due to transparency levels (Kruskal-Wallis rank sum test, $\chi^2[2] = 19.23, p < 0.01$). Post-hoc pairwise compar-

isons using Wilcoxon's rank sum test with Bonferroni correction showed significant differences between the 25% and 50% levels ($p < 0.05$) and between the 25% and 75% levels ($p < 0.01$). The average ratings, grouped according to transparency level and symbol type, are shown in Figure 3.

3.3.2 Foreground interpretation task

For the foreground task, we found main effects for *transparency* ($F[2, 123] = 7.74, p < 0.01, \eta^2 = 0.11$). Tukey's HSD test found a significant difference between transparency levels at 25% and 75% ($p < 0.01$) and a trend towards difference between the 50% and 75% transparency levels ($p = 0.09$). The response times for the foreground reading tasks are shown in Figure 4. As the legend symbols have no bearing on this task, we plotted performance for the two different types of background (rural and urban). As expected, the questions that referred to sliders located over urban areas caused more problems to the participant than those that referred to sliders placed over rural areas. As the transparency level increased, questions took longer to answer. The effect of map region on response time in this case was even stronger than in the background task ($F[1, 20] = 4.16, p < 0.01, \eta^2 = 0.32$) and an interaction between transparency and map region has also been observed ($F[2, 80] = 7.63, p < 0.01, \eta^2 = 0.12$). Tukey's HSD procedure indicated significant differences between urban and rural background ($p < 0.01$) for 25% and 75%.

As in the background case, the difficulty ratings were consistent both with the performance data and with our expectations. Statistically significant effects of transparency on response time were found (Kruskal-Wallis rank sum test, $\chi^2[2] = 23.12, p < 0.01$), but post-hoc pairwise comparisons using Wilcoxon's rank sum test with Bonferroni correction showed significant differences only between the 50% and 75% levels ($p < 0.01$) and between the 25% and 75% levels ($p < 0.01$). No differences were found between 25% and 50% transparency. The average ratings, grouped according to transparency level and map region are shown in Figure 5.
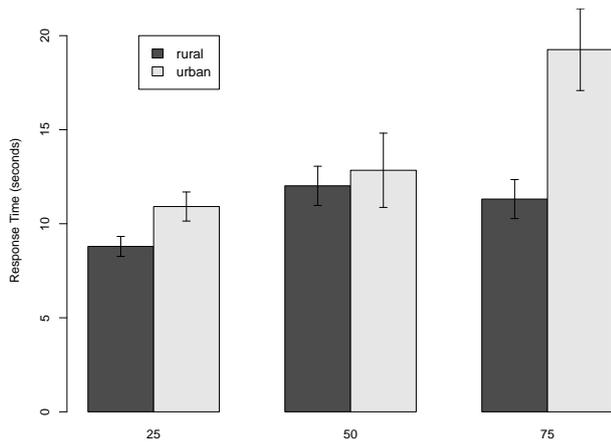
Figure 4: Response times for foreground interface layer tasks, for different foreground transparency levels (25%, 50% and 75%) grouped by background map region over which the tasks were performed.
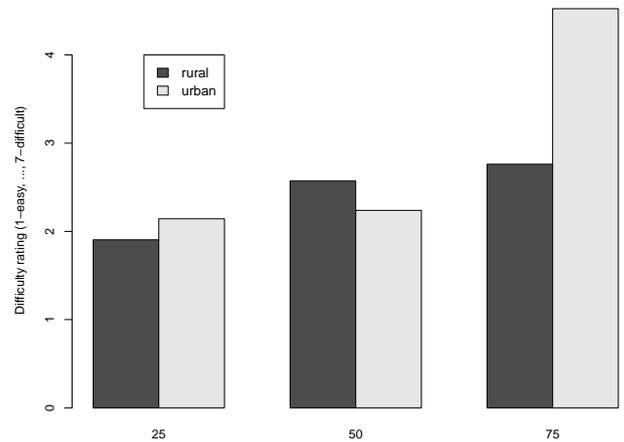

Figure 5: Difficulty ratings for foreground interface layer tasks, for different foreground transparency levels (25%, 50% and 75%) grouped by background map region over which the tasks were performed.

3.3.3 Preferred screen design

As regards user preferences with respect to the three different screens, most users preferred the screen with the 75% transparent foreground layer (54.5% of the participants), while 41% preferred 50% transparency and only 4.5% (a single participant) preferred 25% transparency. This preference for 50% transparency, despite the availability of 75% option, by a large number of participants indicates that they were probably conscious of the need to make selections on the foreground layer easily enough while being able to interpret information shown on the background layer.

4. DISCUSSION

These results show that the main question investigated in this paper, regarding the effect of foreground layer transparency on the user's ability to interpret information shown on the background map layer, has been at least partially settled. Hypothesis H1 has been rejected and significant differences have been shown between 25% and 50% transparency levels. The same data also provide evidence that null hypothesis H3 can be partially rejected. There is strong evidence that reading maps are harder to interpret than seeing maps, but contrary to our expectation this difficulty does not seem to be accentuated by decreasing foreground transparency. Interestingly, however, transparency appears to affect the various symbols types differently. In particular, the questions involving the symbols that denote presence or absence of infected family members (yellow and cyan coloured circles, see Table 1) seem to benefit the most from transparency increases, whereas questions involving the symbols for presence or absence of animals (magenta and green squares) do not seem to be affected. Since continuity (one of the main determinants of transparency perception [1]) is preserved for all cases, it seems that difference of colours could account for the observed performance differences. Relevant in this context is the concept of "colour scissioning"[2], which has been studied in the vision literature

---

[2]One's ability to perceive the colour in the intersecting lay-

for some time [18] and is still an area of ongoing research. Further investigation of colour in semi-transparent map visualization interfaces should take this into consideration.

The rejection of hypothesis H2 in our experiment, which related to the perception of objects on the foreground layer, confirmed the results of Harrison et al. [8]. We also established that for our foreground tasks, based on a map background, 75% is significantly worse than 25% as a level of transparency. However, as we found no significant difference between the 50% and 75% levels, it is reasonable to suggests that 50% transparency may be sufficient for foreground interface layer tasks. This, combined with our results for the background map layer tasks appears to indicate that 50% transparency is the correct design trade-off between preserving the interpretability of map (background) information and preserving the user's ability to manipulate the interface elements on the foreground layer. This conclusion is further corroborated by the answers given by the participants to the final study question, "which of these three settings do you prefer? (25%, 50%, and 75% transparency)", where about 45% of the participants chose the setting displaying the foreground layer at 50% transparency.

Our results also showed that another important factor to consider in the design of transparent interface layers for maps is the type of background. Both foreground and background tasks were significantly harder against urban area bases than against rural area bases, thus rejecting null hypotheses H4a and H4b.

It is conceivable that the type of map visualization interface presented in this paper, in which the user's focus of attention shifts from background to foreground and *vice versa*, could benefit from dynamic adjustments to its transparency levels. In fact, an idea along these lines has been proposed recently by McGookin et al. [17]. It should be noted that, although their approach seemed promising, it fell short due to the fact that changes in transparency levels occurred somewhat unpredictably (from the point of view

---

ers as being composed of the base colour and the overlying transparent colour.

of the user's activity). However, for map visualization tasks such as the epidemiological task that motivated the development of our system, where shifts in focus can often be reliably detected (e.g. by touch event and activity monitoring), allowing the system to automatically vary the transparency of the foreground layer dynamically seems feasible.

While further studies are required, the results above provide an initial set of parameters to guide this and other kinds of design possibilities involving the use of transparency in map visualisation.

5. CONCLUSIONS

We presented a controlled study of the effects of a semi-transparent layer overlaid on an opaque background layer on the interpretability of a map displayed on the latter as the transparency of the former varied. To this end we have investigated different types of map interpretation tasks, different map area types, and their effects on the readability of the foreground layer.

We have found significant effects of transparency level, as well as task and region types, and discussed some possible implications of these results to the design of interactive map visualisation interfaces for exploration of multivariate data through direct manipulation.

In further study we will assess design ideas that involve the use of transparency in dynamic interactive situations. We are specially interested in mobile devices. As the system on which the present study was based is being developed in collaboration with epidemiologists working in the field, we are also considering a longer term observational study to complement the results obtained in controlled laboratory studies so far.

6. ACKNOWLEDGEMENTS

7. REFERENCES

[1] J. Beck and R. Ivry. On the role of figural organization perceptual transparency. *Perception & Psychophysics*, 44(6):585–594, Nov. 1988.

[2] J. Bertin. *Semiology of Graphics: Diagrams, Networks, Maps*. ESRI Press, 2010.

[3] D. B. Carr and L. W. Pickle. *Visualizing Data Patterns with Micromaps*. Chapman and Hall/CRC, 2010.

[4] D. A. Cox, J. S. Chugh, C. Gutwin, and S. Greenberg. The usability of transparent overview layers. In *CHI 98 conference summary on Human factors in computing systems*, pages 301–302, 1998.

[5] M. Elias, J. Elson, D. Fisher, and J. Howell. "Do I live in a flood basin?" synthesizing ten thousand maps. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI'08, pages 255–264, New York, USA, 2008. ACM.

[6] C. Gutwin, J. Dyck, and C. Fedak. The effects of dynamic transparency on targeting performance. In *Proceedings of Graphics Interface*, volume 2003, pages 105–112, 2003.

[7] B. L. Harrison, H. Ishii, K. J. Vicente, and W. A. S. Buxton. Transparent layered user interfaces: an evaluation of a display design to enhance focused and divided attention. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI'95, pages 317–324, New York, USA, 1995. ACM.

[8] B. L. Harrison, G. Kurtenbach, and K. J. Vicente. An experimental evaluation of transparent user interface tools and information content. In *Proceedings of The 8th Annual ACM Symposium on User Interface and Software Technology*, pages 81–90, New York, 1995. ACM.

[9] R. Kasrai, F. A. A. Kingdom, and T. M. Peters. The perception of transparency in medical images. In *Proceedings of The of Medical Image Computing and Computer-Assisted Intervention, MICCAI'99*, volume 1679 of *Lecture Notes in Computer Science*, pages 726–733. Springer, 1999.

[10] M. Kersten, A. Stewart, N. Troje, and R. Ellis. Enhancing depth perception in translucent volumes. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1117–1124, Sept.-Oct. 2006.

[11] T. Koch. *Disease Maps: Epidemics on the Ground*. University of Chicago Press, June 2011.

[12] M.-J. Kraak. Playing with maps explore, discover, learn, categorize, model, analyse, explain, present geographic and non-geographic data. In *Proceedings of the 10th International Conference on Information Visualisation*, IV '06, pages 291–296. IEEE Computer Society, 2006.

[13] H. Lieberman. Powers of ten thousand: navigating in large information spaces. In *Proceedings of The 7th Annual ACM Symposium on User Interface Software and Technology*, UIST '94, pages 15–16. ACM, 1994.

[14] S. Luz, M. Masoodian, and M. Cesario. Disease surveillance and patient care in remote regions: an exploratory study of collaboration among healthcare professionals in amazonia. *Behaviour & Information Technology*, 2013. Published online Dec 2013.

[15] A. M. MacEachren. *How Maps Work: Representation, Visualization, and Design*. The Guilford Press, 1995.

[16] A. M. MacEachren and M.-J. Kraak. Research challenges in geovisualization. *Cartography and Geographic Information Science*, 28(1):3–12, 2001.

[17] D. McGookin, I. Herteleer, and S. Brewster. Transparency in mobile navigation. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems*, pages 1903–1908, New York, 2011. ACM.

[18] F. Metelli. The perception of transparency. *Scientific American*, 230(4):90–98, 1974.

[19] E. R. Tufte. *Envisioning information*. Graphics Press, 1990.