Single Complex Glyphs Versus Multiple Simple Glyphs

Beth Yost, Chris North
Department of Computer Science
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061 USA
{beyost, north}@vt.edu

ABSTRACT
Designers of information visualization systems have the choice to present information in a single integrated view or in multiple views. In practice, there is a continuum between the two strategies and designers must decide how much of each strategy to apply. Although high-level design guidelines (heuristics) are available, there are few low-level perceptual design guidelines for making this decision. We performed a controlled experiment with one, two, and four views to evaluate the strengths and weaknesses of these strategies on target detection and trend finding tasks in the context of multidimensional glyphs overlaid onto geographic maps. Results from the target detection tasks suggest that visual encoding is a more important factor when detecting a single attribute than the number of views. Additionally, for detecting two attributes, the trend indicates that reusing the most perceptually salient visual feature in multiple views provides faster performance than an integrated view that must map one of the attributes to a less salient feature.

Author Keywords
Information visualization, multiple views, evaluation.

ACM Classification Keywords
H5.2 [User Interfaces]: Evaluation/methodology.

INTRODUCTION
In the design of information visualizations, one design dimension that designers must wrestle with is the continuum between two opposing design strategies: integrated views and multiple views. The integrated view strategy seeks to represent all of the information in a single integrated view. It emphasizes the advantages of layering information and integrating more information into one view. The multiple view strategy splits information into multiple separate linked views [6]. It emphasizes the advantages of segmenting complex information into simpler parts. In practice, designers must decide how much of each strategy to apply for a given dataset and user task.

Baldonado et al. [1] provide some heuristics for deciding when it is advantageous to split data between multiple views instead of integrating the data into a single complex view. In [8] some of the cognitive issues involved are empirically evaluated. Burns has studied both multiple view and integrated view visualizations presented simultaneously and sequentially (although she uses different terminology) [2] in the specific domain of a nuclear power plant operator. Trafton et al. have studied multiple visualizations presented to Navy meteorologists sequentially [9]. Research has also been conducted on user performance with different methods of linking views [5], and on the effectiveness of different combinations of views for the exploration of multivariate health data [3].

Most of these studies are user performance comparisons in specific domains and for particular types of users either with multiple views or with integrated views – but not a comparison between both strategies. Some advocate the integrated views approach [4] while others advocate multiple views [7]. The guidelines that do exist ([1]) are not currently based on empirical evaluation. This opens the door to many research questions about visualization design, and in particular, about how much of each design strategy is appropriate for particular tasks.

METHOD
The goal of this experiment is to examine the tradeoff between integrated and multiple views. The focus is on visualizing geospatially-referenced multi-dimensional data points, as is common in GIS, using simple glyph-based encoding (as in [4]). Data points are mapped to visual glyphs and overlaid on a geographic map. This represents a common visualization problem in which users must relate spatial and multi-dimensional data. This study does not consider other data representations (such as in [3][5]).

The primary tradeoff is as follows: In a larger integrated view, multiple data attributes are encoded using complex glyphs with multiple visual features (such as color, size, and orientation) which could potentially interfere with each other. In multiple views, each data attribute is represented in a separate smaller view, using simple glyphs with one most-salient visual feature (e.g. color), as in Tufte's "small multiples" [10]. The complexity tradeoff is between the number of visual features of glyphs and the number of views. Presumably, both should be minimized to reduce perceptual and cognitive complexity. What is the best approach for different types of tasks?

CHI 2005, April 2–7, 2005, Portland, Oregon, USA.
ACM 1-59593-002-7/05/0004.

Experiment Design

This experiment is a 3x2x4 design with 3 independent variables: number of views, task attributes' visual encoding, and task type. It is a mixed design. Number of views is between-subjects, while the others are within-subject.

The number-of-views variable had 3 conditions: one, two, and four views were compared (Figures 1, 2, and 3). The two-view case is an intermediate solution between integrated and multiple views. In each condition, four abstract data attributes were shown – attributes A, B, C, and D. The actual data used was discrete and a modified version of data obtained from the US Census. Data attributes were distributed equally among views, and each visual encoding was used only once per view as shown in Table 1. The visual mappings used for representing these attributes have already been proven effective in [4]. Attribute A was considered the most important and mapped to the best encoding (color [4]), and so on with D the least important. For simplicity in reporting results, we refer to A as having a 'better encoding' than B. In all conditions, total screen space used was held constant.

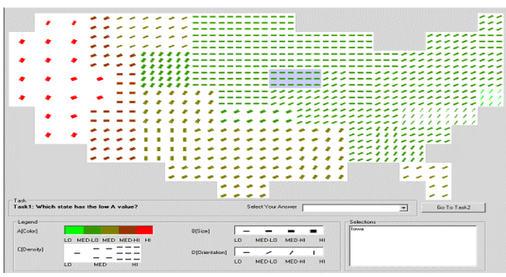

Figure 1. Integrated, 1 view. Data attributes mapped to color, size, density, and orientation.

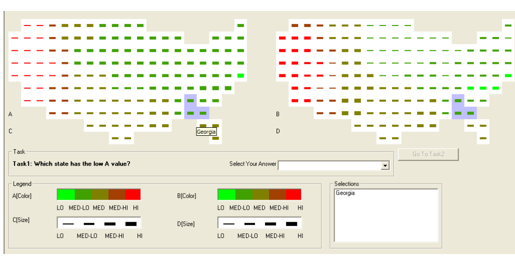

Figure 2. Dual, 2 views. Color and size used in both.

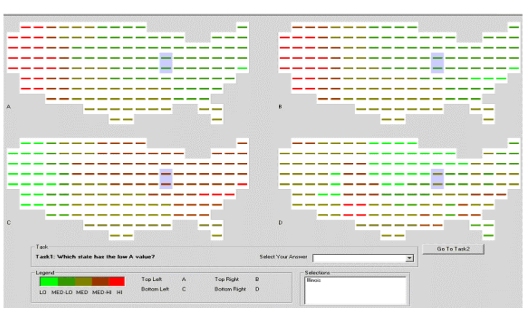

Figure 3. Multiple, 4 views. Color used in all four.

# Views		Data Attribute			
		A	**B**	**C**	**D**
	1	Color	Size	Density	Orientation
	2	Color (left)	Color (right)	Size (left)	Size (right)
	4	Color (top left)	Color (top right)	Color (low left)	Color (low right)

Table 1. Visual encodings used in each condition of # of views.

	Most important attribute, Best visual encoding	Less important attribute, Not best visual encoding
Detect:One attribute	Which state has the lowest A value?	Which state has the highest D value?
Detect: Two attributes	Which state has the medium low A value and the low B value?	Which state has the medium low A value and the medium C value?
Trend: One attribute	What's the trend from West to East in terms of A value?	What's the trend from North to South in terms of D value?
Trend: Two attributes	What's the relationship between A and B?	What's the relationship between A and C?

Table 2. Four tasks and 2 conditions for visual encoding.

Four different tasks were completed: target detection with 1 attribute and with 2 attributes, and trend finding with 1 attribute and with 2 attributes. These tasks can be seen in the rows of Table 2.

Each task also has 2 variations (visual-encoding variable, columns of Table 2): the first asks about the most important attributes (A and B), while the second includes a less important attribute (C or D). The purpose of this variable is to explore the effect of the tradeoff between glyphs features and views. The most likely affected data attributes are those that are of less importance, and hence mapped to less effective visual encodings in the case of integrated views.

The experiment measured three dependent variables. First, user performance time was measured for each task. Second, answer correctness was recorded. The last measure was taken on a post-questionnaire and was a subjective measure of how well the participant believed the interface supported the different types of tasks.

The fifty-seven participants in this study were engineering students from a large public university. Participants performed practice tasks for detecting one and two targets. During the practice, the administrator explained the various features in the interface. There was no time limit for each task. At the end of the experiment, participants were given a post-questionnaire to rate the interface.

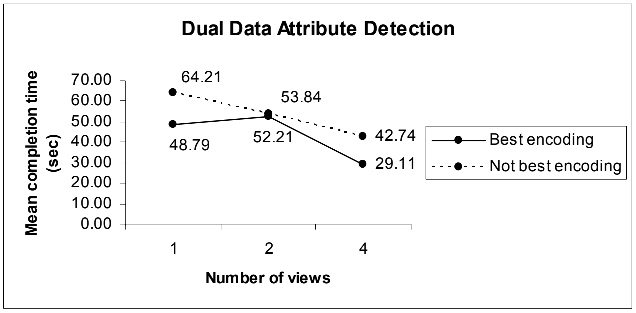

Figure 4. Completion times for detecting two attributes.

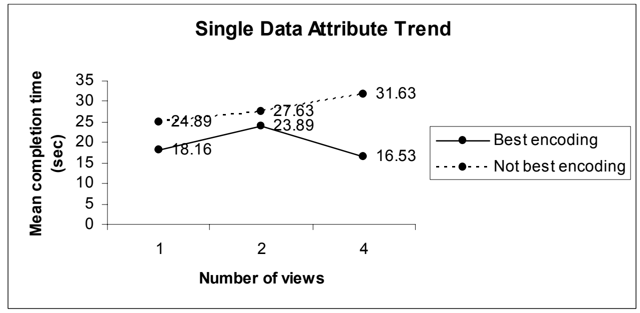

Figure 5. Completion times for one-attribute trend task.

RESULTS

Since it is not of interest to compare between the 4 tasks, each is treated as a separate 3x2 analysis. Two-way ANOVAs were performed for each task (each row in Table 2) on the other 2 factors: number of views, visual encoding.

Single Attribute Detection Tasks

For the tasks that involved detecting one data attribute there was a significant effect of number of views, $F(2, 54) = 12.06$, $p < .01$, visual encoding, $F(1, 54) = 14.06$, $p < .01$, and an interaction effect between them, $F(2, 54) = 11.85$, $p < .01$. Detecting the attributes that had the better visual encoding was significantly faster.

Tukey's HSD for the number of views indicated that participants that used one view took significantly less time than those using two views, and those using four views took significantly less time than those using two views (both $p < .01$). There was not a significant difference between one and four views.

For the interaction effect between number of views and visual encoding, task completion times were significantly faster when using two views with the best visual encoding than when using two views with the less perceptually salient encoding ($p < .01$). When detecting a single data attribute with the less perceptually salient visual encoding, task completion times were significantly faster when using one view compared to two views, and also when using four views compared to two views ($p < .01$).

There was no significant difference in correctness. In terms of satisfaction, anova indicated that there was a significant effect of the number of views, $F(2, 54) = 3.57$, $p < .05$ for these tasks. Tukey's HSD for the number of views indicated that participants that used two views were significantly less satisfied than those using four views ($p < .05$).

Dual Attribute Detection Tasks

For detecting two data attributes there was a significant effect of visual encoding, $F(1, 54) = 4.83$, $p < .05$. Detecting the two attributes with better visual encodings took significantly less time. Figure 4 shows the mean completion times for tasks that involved detecting two attributes. There was not a significant effect for the number of views or their interaction. There was no significant difference in terms of correctness or satisfaction.

Trend Finding Tasks

For the tasks that involved finding a trend for a single attribute there was not a significant effect by number of views, but there was by visual encoding $F(1, 54) = 23.94$, $p < .01$ and for their interaction $F(2, 54) = 3.81$, $p < .05$. Finding the trend for the attribute that was less perceptually salient took significantly longer. Tukey's HSD for the interaction indicated that, when using one view, finding the trend with the best visual encoding was significantly faster than finding the trend with a less perceptually salient visual encoding. The same was true for four views ($p < .05$). The mean completion times for this task can be seen in Figure 5. There was no statistically significant difference between views in terms of correctness or satisfaction for this task.

For the task that involved finding a relationship trend between two data attributes, there was no significant difference in terms of time, correctness, or satisfaction.

DISCUSSION

The decision of how many views to use when designing a visualization involves tradeoffs. Using more views increases the complexity of the interface, but allows more perceptually salient visual encodings to be reused. When using a single integrated view the problem may also arise that all reasonable visual encodings have been exhausted.

Target Detection Tasks

The results of both the single and dual attribute detection tasks indicate that it is faster to find the attribute with the best visual encoding. What is initially surprising is that, in the single attribute detection task, dual views seemed to perform worse than integrated and multiple. This suggests that the visual encoding of size may have been harder to detect than either color or, in particular, orientation. It is also possible that the dual view condition performed worse because it was harder for participants to understand the legend. In a single view, all attributes are represented by different encodings in a single map. In multiple views, every map represented a different attribute. The dual view was a combination of both strategies and hence required users to decode both methods. Particularly in the dual view situation, legends must be carefully designed.

There was no significant difference when color was used in each, but there was a significant difference when different visual encodings were used for the same task. Therefore, it appears that visual encoding has more of an impact on user performance than the number of views and any visual interference due to extra visual encodings. The lesson learned from this is that it is faster to use the best visual encoding for the data attributes of interest. Either the visualization can be designed with the ability to change the visual encodings of attributes as the user's interest changes, or multiple views can be used because of the ability to reuse the best visual encoding for several attributes. Color encoding was always the best in this study.

Trend Finding Tasks

The results of the trend finding task on one attribute indicated that using a better visual encoding was faster. Although this is expected when using a single view with different encodings, it is not when using multiple views. Aside from the trend direction and data, the only difference was the location of the target view in the 4-view grid. It is possible that this task took longer because participants became confused with which map to consider. This was only true when the trend for one attribute needed to be found. The lesson learned from this is that if multiple views are used there is the potential for confusion as more views are added, particularly when the task becomes more cognitively demanding. To avoid this problem the interface could be designed so that views can be rearranged or highlighted. At a minimum, the interface should be clear regarding which attributes are represented in each view.

Exploratory Results

In the dual attribute detection task, the best visual encoding condition involved finding color and size in a single view (integrated), or color in two different views (dual and multiple). The trend in the results indicates that using multiple views took less time than using the dual view interface. At first it seems these should be close to equal considering they both involve finding color in two different views. The difference between these is that in the dual view situation there are extra visual encodings, and in the multiple view situation there are extra views. Therefore, it appears that it is easier to identify colors in two views if no extra visual encodings are present in those same views.

Finding the attributes with the less perceptually salient encodings involved finding color and density in a single view (integrated), color and size in a single view (dual), or color in two different views. The trend indicates that multiple was faster than dual which was faster than integrated (Figure 4). This means that using color in two different views was faster than using either color and orientation or color and size in a single view. It appears that being able to use the two best visual encodings in different views was faster than showing both combined when one of the data attributes had a visual encoding that was less salient. These data trends warrant further research.

CONCLUSION

The main finding is that visual encoding is the most important factor, regardless of number of views. Furthermore, the trend indicates that it is faster to reuse the most perceptually salient visual feature in multiple views than it is to use an integrated view that requires a less perceptually salient feature in complex glyphs, even for two-attribute tasks. Since integrated views never out-performed multiple views, it appears that the multiple-views strategy is the safest decision to guarantee good performance via effective visual encodings. Alternatively, users should be able to change the visual encodings of attributes as their attribute importance changes. Results also suggest that legends should be carefully designed.

Future work includes increasing the number of data attributes, views, visual encodings, and task complexity, as well as exploring alternative data representations and interactivity. The distribution of attributes among views must be studied to determine the effect of repeat encodings. This could lead to a stronger theory for visualization design that encompasses multiple-view strategies.

ACKNOWLEDGEMENTS

This research is partially supported by NSF grant #CNS-04-23611, the National Geospatial-Intelligence Agency ARDA GI2Vis program, and the US Bureau of the Census.

REFERENCES

1. Baldonado, M.Q., Woodruff A., and Kuchinsky A. Guidelines for Using Multiple Views in Information Visualization. In *Proc. AVI2000*, (2000), 110-119.
2. Burns, C.M. Putting It All Together: Improving Display Integration in Ecological Displays. *Human Factors 42*, 2 (2000), 226-241.
3. Edsall, R.M., MacEachren A.M., and Pickle L. Case Study: Design and Assessment of an Enhanced Geographic Information System for Exploration of Multivariate Health Statistics. In *Proc. IEEE Symposium on Information Visualization*, IEEE (2001).
4. Healey, C.G., St. Amant, R., and Elhaddad, M. ViA: A Perceptual Visualization Assistant. In *28th Applied Imagery Pattern Recognition Workshop*, (1999), 1-11.
5. Li, Q., and North, C. Empirical Comparison of Dynamic Query Sliders and Brushing Histograms. In *Proc. IEEE Symposium on Information Visualization*, IEEE (2003).
6. North, C. Multiple Views and Tight Coupling in Visualization: A Language, Taxonomy, and System. In *CSREA CISST 2001 Workshop on Fundamental Issues in Vis.*, (2001), 626-632.
7. Roberts, J.C. On Encouraging Multiple Views for Visualization. In *Proc. IV'98*, (1998), 8-14.
8. Ryu, Y.S., Yost, B., Convertino, G., Chen, J., and North, C. Exploring Cognitive Strategies for Integrating Multiple-View Visualizations. In *Proc. HFES 2003*.
9. Trafton, J.G., Kirschenbaum, S.S., Tsu, T.L., Miyamoto, R.T., Ballas, J.A., Raymond, P.D. Turning Pictures Into Numbers: Extracting and Generating Information from Complex Visualizations. *International Journal of Human-Computer Studies 53*, 5 (2000), 827-850.
10. Tufte, E., *Envisioning Information*, Graphics Press, 1990.