

How Capacity Limits of Attention Influence Information Visualization Effectiveness

Steve Haroz and David Whitney

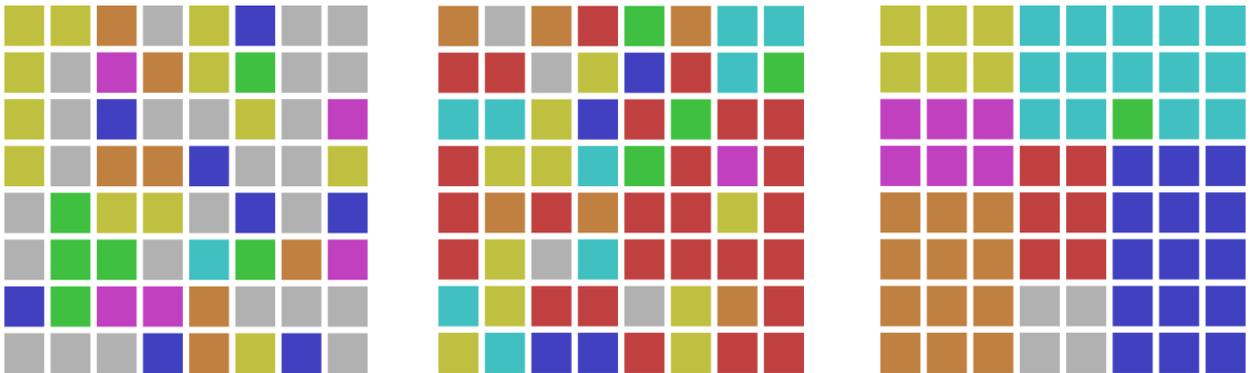


Fig. 1. These images each have one colored square that is unique within that image. How long does it take you to find each? How many color categories are there in each panel? Why does grouping make both tasks substantially easier?

Abstract—In this paper, we explore how the capacity limits of attention influence the effectiveness of information visualizations. We conducted a series of experiments to test how visual feature type (color vs. motion), layout, and variety of visual elements impacted user performance. The experiments tested users' abilities to (1) determine if a specified target is on the screen, (2) detect an oddball, deviant target, different from the other visible objects, and (3) gain a qualitative overview by judging the number of unique categories on the screen. Our results show that the severe capacity limits of attention strongly modulate the effectiveness of information visualizations, particularly the ability to detect unexpected information. Keeping in mind these capacity limits, we conclude with a set of design guidelines which depend on a visualization's intended use.

Index Terms—Perception, attention, color, motion, user study, nominal axis, layout, goal-oriented design.

1 INTRODUCTION

An information visualization designer aims to present the maximum amount of data without overwhelming the user with complexity and information-overload. The components arranged to form a GUI or visualization are visual features – the properties of any image that the brain is capable of encoding and integrating into a coherent percept [1]. Examples that apply to visualization include position, color, size, orientation, texture [2], and motion [3–5]. The designer's role is to effectively associate these visual features with corresponding dimensions or categories in the underlying data [6].

Unfortunately, the speed and capacity of human attention for these visual features are severely limited [7], [8] and these limits may influence the effectiveness of information visualizations. Exceeding the limits of visual attention markedly impairs both the accuracy and timing of one's response to a visual scene (Fig. 1). This consequence may seem intuitive (e.g., the benefit of grouping in Fig. 1 might seem obvious), however visualizations often violate or ignore this intuition in part because it has not been formalized or empirically tested in visualization-relevant tasks. Characterizing and

measuring these limitations in the context of data visualization is therefore fundamental and necessary to achieve the goal of conveying information via the human visual system.

Using a visualization should provide information faster or more broadly compared with serially inspecting the raw data in the form of a table or database [9]. A perceptual hindrance that restricts a user to serially inspecting each visualized element rather than enabling a rapid summary perception of the whole scene would make the visualization hardly better than a simple table. It is therefore critical to understand how the capacity limits of attention impact various visualization tasks.

To study the effect of these limits, we tested the effect of two types of arrangements or layouts on user performance in visual search and subitizing (or rapid counting [10]) experiments. We also examined how performance in these experiments is influenced by user goals and the variety of a visualization's visual features. To measure maximum performance, the experiments all had subjects fully attending to one task as opposed to dividing attention with a peripheral display [11].

We ran three experiments with the similar stimuli that used either colored or moving features. Each experiment's task, however, corresponded to a different, commonly performed visualization task:

- Detect a unique target with a known appearance (e.g., find the red object)
- Detect a unique target with an unknown appearance (e.g., find a unique or oddball target)
- Determine and compare the number of visual categories (e.g. determine extent of heterogeneity or consistency)

The latter two tasks are of particular interest, as performing them by browsing a table or running a database query is difficult. For example, finding the hour of the day that you receive the most or

• Steve Haroz is with University of California, Davis, e-mail: sharoz@ucdavis.edu, website: <http://stevharoz.com>

• David Whitney is with University of California, Berkeley, e-mail: dwhitney@berkeley.edu, website: <http://whitneylab.berkeley.edu>

Manuscript received 31 March 2012; accepted 1 August 2012; posted online 14 October 2012; mailed on 5 October 2012.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

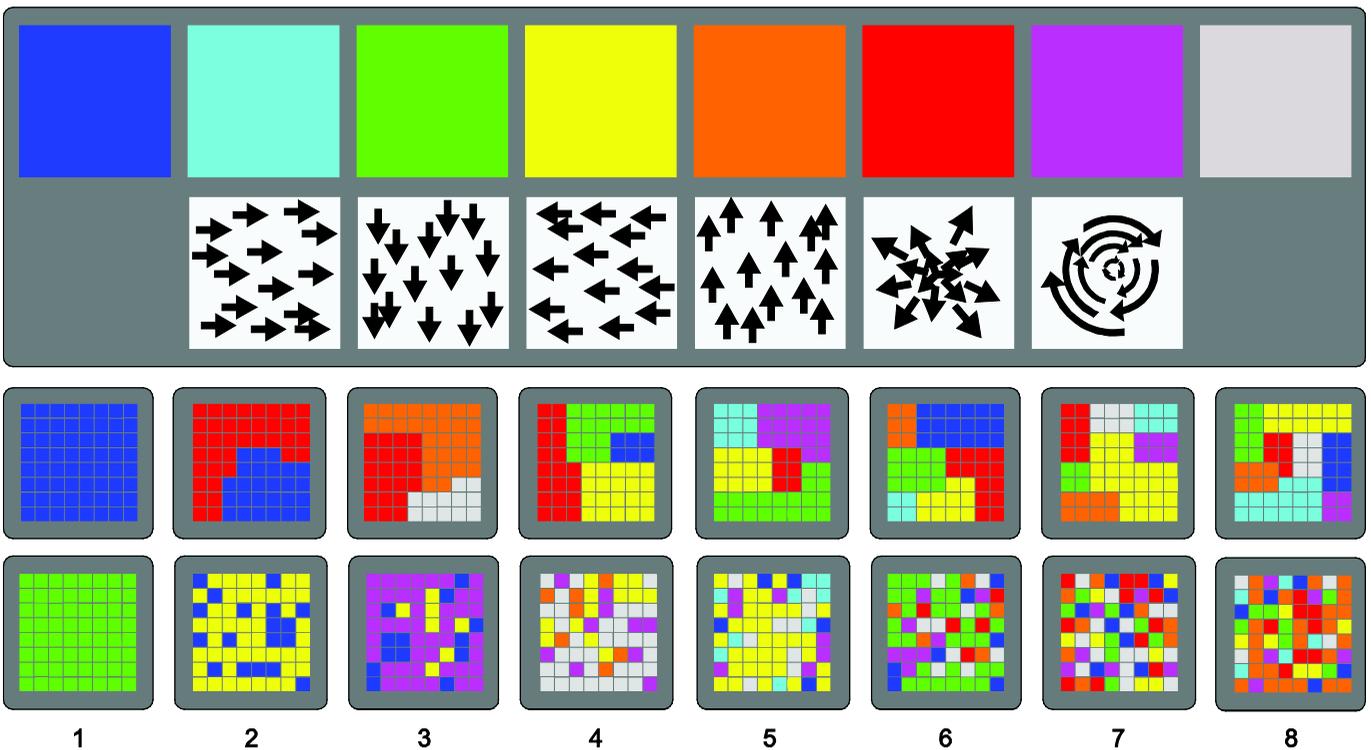


Fig. 2. (Top) All of the colors and motion types used in the stimuli. The saturation is enhanced in the figures, as the actual colors were more isoluminant. (Bottom) A comparison of the grouped and random layouts for differing variety amounts. Each variety amount (except 1) had 16 possible grouped layouts (manually created by the authors). Each variety had a different number of squares in any given stimulus.

fewest emails would be nontrivial without some sort of graphical display. Such tasks which assume no prior knowledge of where to look or what to look for are where information visualization excels.

One of our primary goals was to quantify the influence of a grouped arrangement compared with a random arrangement (e.g., Fig. 1) on visual search and subitizing tasks. We hypothesized that the effect would vary by task and would modulate the impact of capacity limits, so each experiment separately tested both layouts. To show that these effects were consistent across visual feature, we tested both color and motion. Several studies have investigated the impact of variance and grouping on visual search. Duncan and Humphreys showed that visual search of a known target (or one of two known targets) can be impacted by variance among even a small number of objects [12]. Treisman found that increasing the attentional demand by searching for a conjunction of visual features increases the adverse impact of more groups [13]. In the context of HCI and visualization, the spatial and sizing consequences of grouping have been proposed as models for design and evaluation [14], [15]. Furthermore, Tatu et al [16] examined how grouping in scatter plots influences preference (grouped displays are preferred), but preference and performance (an operational measure of visualization effectiveness) are not necessarily equivalent. We aim to extend these investigations of search, attention, and capacity and present our findings in the context of information visualization.

For each experiment and each block, we tested how the variety of a feature (color or motion) affected performance. This property corresponds to the number of categories or discrete steps of a nominal axis that are displayed in a visualization. While showing the largest amount of detail or dynamic range for each data axis is ideal, we tested whether adding more detail can actually reduce someone's ability to utilize any of the information.

2 EXPERIMENT STIMULI

All three experiments used the same stimuli. They consisted of an 8 x 8 grid of 64 squares centered on a gray background (similar to Fig. 1). The number of squares remained constant for all of the experiments to emulate simple visualizations with the same number

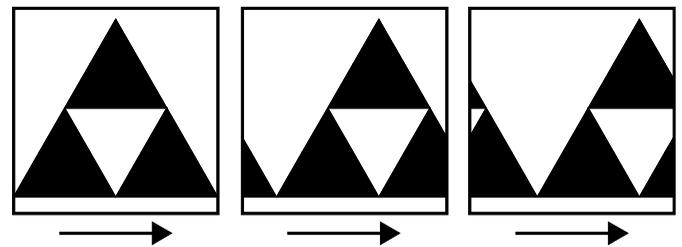


Fig. 3. A demonstration of the textured motion used in the experiments. The texture moves, but the square's position never changes. All motions cycled at 1 Hz.

of data elements but differing arrangements (correlation with the spatial axes or grouping algorithm) and amounts of variety (categorical detail).

For the color trials, each square had one of the eight highly discriminable colors in Fig. 2. The large number of easily discriminable colors is what makes color such a frequent and effective tool for visualization [17], [18]. For the motion trials, each square acted as an aperture onto a monochromatic tiled texture (Fig. 3) which moved using one of the motions in Fig. 2. The number of motions was limited to six to ensure easy discrimination of each type of motion. The texture was the same for all squares but varied between trials. All of the textures had the same number of black and white pixels, so overall brightness did not vary.

As the top of Fig. 2 shows, the stimuli had a limited number of colors or motions. The subset of colors or motions was randomized for each trial. We refer to the number of colors or motions in each stimulus (not including a search target) as the amount of 'variety'. Furthermore, the experiments were divided into grouped layouts and random layouts (bottom of Fig. 2). For the grouped layouts, all squares of a particular color or motion were grouped into a single cluster. The sizes of the groups varied within a given display, so the size of an individual group provided no information about the number of groups. The random layouts assigned a color or motion to each square using a $1/n$ random distribution. For both layouts each

color or motion was present on at least four squares.

All of the experiments were divided into four blocks:

- Color – Grouped
- Color – Random
- Motion – Grouped
- Motion – Random

The order of these blocks was balanced between subjects such that each block went first and last at least once.

The experiment application was written in XAML and C# for the Windows Presentation Framework. The motion component was an HLSL shader that independently applied a transform in texture-space to each textured square. The experiments were conducted on a 24 inch iMac (using the built-in LCD monitor) running Windows 7 via Boot Camp.

The subjects performed the experiment in a dark room to avoid external visual distractions. They maintained a distance of 22.4 inches (57 cm) from the screen via a chinrest. The stimulus was 12.7 inches (32 cm) tall and wide. It consequently encompassed 32.3° of visual angle, while the squares were each 3.7° due to the small gaps between them.

For all the results with response time (RT) analyses, trials with an absent target (50%) or that had an incorrect response were not used.

3 VISUAL SEARCH: FIND A KNOWN TARGET

The first experiment was the simplest and served primarily as a control. The aim was to simulate the task of finding a known target. Realistic examples include finding a Green Party region on an election map, finding a large file in a directory treemap, or finding rain in a weather visualization. In these scenarios, the user knows the appearance of the target (perhaps via a legend) and is searching for its presence.

3.1 Methods

Five subjects participated in this experiment. Two were female. All were either graduate students in psychology or computer science or trained university staff.

As Fig. 4 shows, each trial began by displaying a square with the target visual feature in the center of the screen. A progress bar for the elapsed number of trials was also displayed. A one second pause (a gray screen) was then displayed to prevent any aftereffects or apparent motion. Then the stimulus was presented. Using the keyboard, the subject responded whether the target was *present* or *absent*. Answers and response times were recorded.

For each block, each variety count (number of colors/motions in the stimulus) had 40 trials. In half of the trials, a random square was replaced by the target, whereas the other trials had no target. Each of the experiment's color blocks had 280 ($7 * 40$) trials, and each motion block had 200 ($5 * 40$) trials for a total of 960 trials per subject (4,800 total trials). The trials were randomly ordered within each block, which began with five practice trials.

3.2 Results

The results in Fig. 5 show several trends (ANOVA results included):

Accuracy was over 95% for all blocks and did not diminish with higher variety counts.

Visual feature: Color had a trend for lower RTs compared with motion for grouped – $F(1,3) = 921, p < 0.0001$ – and random – $F(1,3) = 41, p < 0.01$. This effect was not unexpected and reaffirms that visual feature can impact user performance.

Layout: The grouped layout performed slightly better than the random layout for color – $F(1,4) = 65, p < 0.001$ – and motion – $F(1,3) = 19, p < 0.05$. Size pop-out (when a distinct object perceptually stands out from its surroundings [1]) was possibly the reason because the grouped layout had no other groups as small as the target (1 square). For the grouped layout, detecting if a square was unique amongst its immediate neighbors could confirm it as a

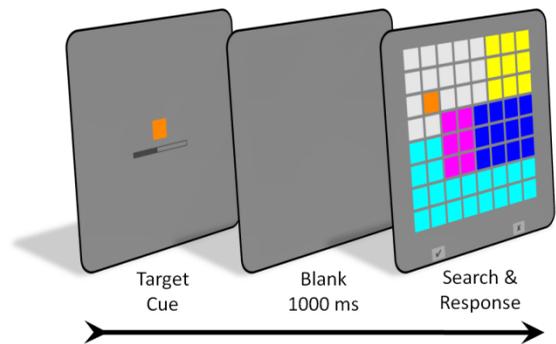


Fig. 4. The procedure for the visual search experiments. The target and a progress bar were shown to the subject (the oddball experiment only had a progress bar). A blank gray screen was presented to avoid any aftereffects or apparent motion. Finally the stimulus was displayed while awaiting the subject's response.

target without needing to check the entire screen for confirmation.

Variety: For color and grouped motion, the variety count had almost no impact on performance. The random motion block's RTs appeared to increase until 3 varieties and then plateau thereafter like the other blocks.

Subject: Though between-subject variability was significant, all subjects had the same trend of results (Fig. 5 error bars).

A purely horizontal slope implies that the visual system can detect the presence of a target in parallel irrespective of the number of varieties. Duncan and Humphreys' study [12] found a very small performance cost of a few milliseconds per additional category. Our results suggest that this cost remains minimal for a large number of objects.

4 VISUAL SEARCH: FIND THE ODDBALL

This variant of the visual search experiment served to test users' ability to detect something odd or out of place. It is analogous to finding a unique element in a visualization without knowing to look for it. Examples include noticing that one file's icon is different compared with its neighbors in a file manager or detecting an oddly behaving port in a network security visualization. This type of scenario is where visualizations can be most useful because if the user knows what to search for, a simple database query may yield faster and more accurate results. Detecting a unique element—an oddball, however, is nontrivial to query.

4.1 Methods

Five subjects participated in this experiment. Four were female. All were either graduate or post-graduate students in psychology or trained university staff.

The experiment was identical to the known-target experiment shown in Fig. 4, but the target was not shown (only the progress bar was visible). Subjects were told simply to detect if a unique square was present. As in the known-target experiment, each subject performed 960 trials (4,800 total trials).

4.2 Results

The results in Fig. 6 show several trends (ANOVA results included):

Accuracy: In contrast to the known-target search experiment, accuracy in finding the oddball was significantly affected by the amount of variety. Although the grouped layouts showed little performance degradation with increased variation, the random layouts had a marked decline as variation increased.

Visual feature: Color again had consistently lower RTs than motion for grouped – $F(1,4) = 120, p < 0.0001$ – and random – $F(1,4) = 101, p < 0.0001$. Nevertheless, they both showed similar trends despite the differing scales.

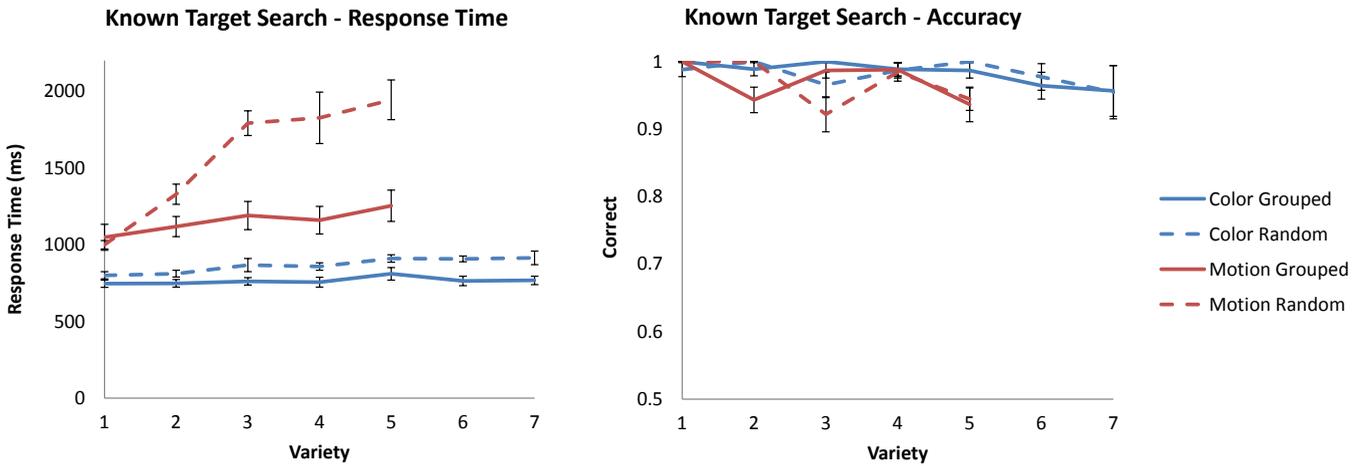


Fig. 5. Experiment 1 results: visual search for a known target. (Left) The RTs of the visual search experiment as a function of the variety of features (e.g., the number of colors or motion directions visible). Only correct responses for present targets were used in the calculation. (Right) The accuracy was consistently above 95%. Error bars show the inter-subject standard deviation.

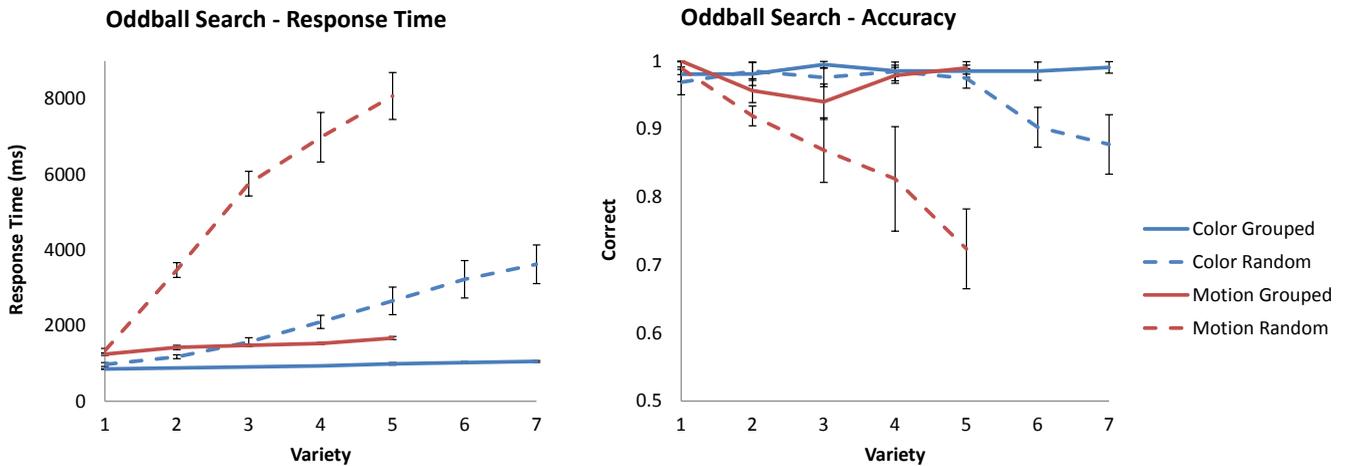


Fig. 6. Experiment 2 results: visual search for an oddball target. (Left) RTs plotted as a function of the variety of features. Only correct responses for present targets were used in the calculation. Note that the variety of features has little influence on the RTs for grouped layouts, while the RTs for random layouts clearly grow with more variety. (Right) Accuracy for the grouped layouts was consistently over 90%, whereas the accuracy for random layouts dropped significantly as variety increased. Error bars show the inter-subject standard deviation.

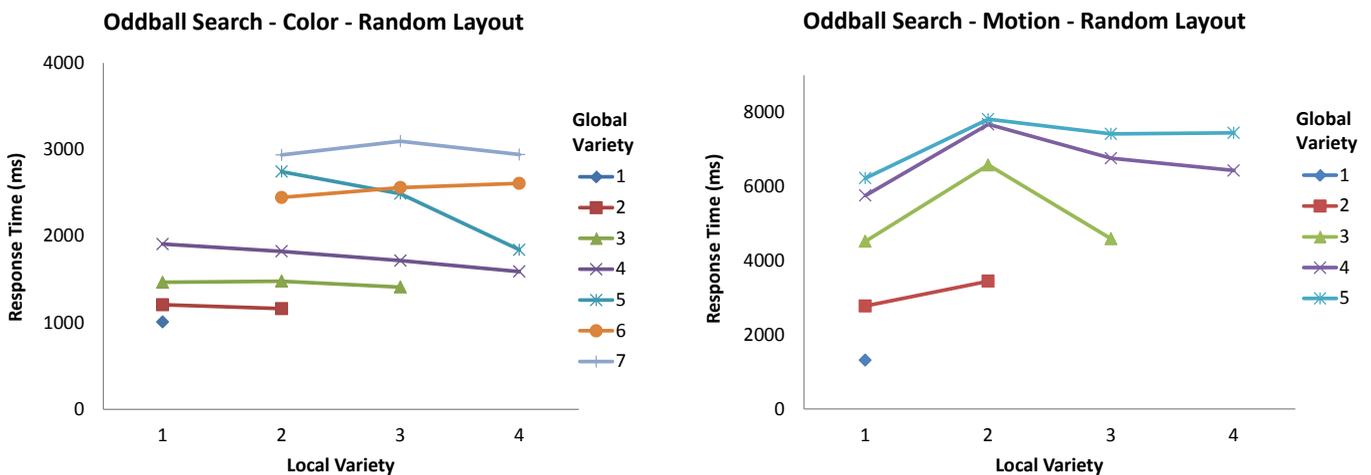


Fig. 7. Local vs global variety. These plots show the effect of local variety on performance under different amounts of global variety. Local variety is the number of colors or motions immediately adjacent to the target, whereas global variety is the total number of colors or motions on the screen. The horizontal slopes show the lack of impact of local variety on user performance.

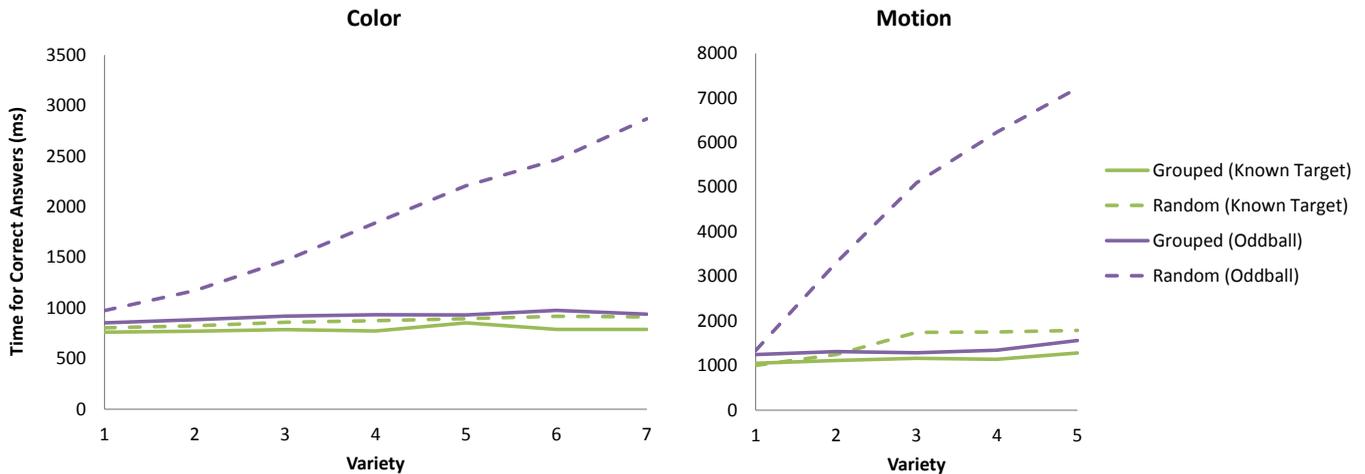


Fig. 8. The figure shows the response times for the visual search experiments with known and oddball targets and groups them by visual feature. In spite of their different scales, the same trend is clear – variety of visual features has little influence on RT when subjects know the target in advance or have a grouped layout. That is, targets are relatively easy to find when they are predefined or when the distractors are grouped. Finding an unknown target in a random layout, on the other hand, becomes significantly more difficult with increased variety. From this information, we can infer that tasks requiring users to find oddball targets among ungrouped layouts (a likely occurrence with scatter plots, treemaps, or connectivity graphs) would make for an ineffective visualization design.

Layout: The grouped layout performed significantly better in RT and accuracy for both color – $F(1,6) = 109, p < 0.0001$ – and motion – $F(1,4) = 117, p < 0.0001$. In the grouped blocks, subjects were capable of detecting the target without influence from the amount of variety. Increasing the variety in the random layouts markedly impaired performance.

Variety: For the grouped blocks, the variety amount did not affect RT or accuracy. Conversely, the random blocks showed a clear dependence on the number of visual features used – $F(1,4) = 22, p < 0.05$

Subject: Though the between-subject variability was significant, all subjects had the same trend of results (Fig. 6 error bars).

Local vs. global variety: We sought to examine whether the performance difference was caused by variety in the whole scene or immediately surrounding the target. To calculate local variety, we counted the number of colors or motions in the eight squares surrounding the target. Fig. 7 shows the RTs for the random layouts plotted by local and global variety. The grouped layouts showed no effect from either type of variety.

The strong distinction between Experiments 1 and 2, searching for a known target versus an oddball target, is apparent in Fig. 8. Corroborating the results of Nothdurft [19], color and motion show the same pattern. The variety of visual features does not strongly impact performance when the search target is known, but has a dramatic effect when the search target is not known (is an oddball). Interestingly, grouping the features mostly compensates for this difference, creating a large performance gap between grouped and random layouts for the oddball search. The implication is that subjects are switching from a parallel process in the known-target and grouped odd-ball tasks to a serial process dependent on the number of features. A possible explanation is that the subjects were performing a series of Boolean map filters [20] to detect global uniqueness only for the random oddball search. How much a target *popped-out* from its neighbors (due to lack of surrounding variety) had no impact on performance. This process would be repeated until the search yields no similar squares. These Boolean map filters would be applied sequentially for each variety until a unique target was found.

5 SUBITIZING: LIMITED CAPACITY

The goal of this experiment was to understand how visual complexity affects a user’s ability to grasp aspects of the global structure of a display. For example, users may need to extract

information about the gist of a visualization (was the market, on average, up or down in Fig. 11), or they may need to estimate the variety of types of information (how many unique file types are in a tree map? Or, how much variance is there in a map of the market—was the market mixed today, or was it homogeneously down?). The task in our experiment was to subitize (or rapidly count [10]) the relative number of categories in a pair of stimuli.

Understanding the overall structure and number of categories in a dataset is yet another task which would be difficult without a visual aid (e.g., finding how many file types are in a directory is trivial with a nested tree map visualization). Yet even these aids are not always particularly helpful, as many have experienced struggling to understand the overall message of a display with a confusing layout or too many visual categories. This struggle highlights the fact that it is unknown how overall qualitative comprehension is affected by layout or visual feature variety. This experiment sought to provide some insight on the matter. Practical examples are available however. The NewsMap (<http://newsmap.jp>) in Fig. 12 shows an example of similar information being either easy or difficult to discern depending on the layout. When the grouping deteriorates, even a basic gist (or ensemble percept [21], [22]) of the most significant news category becomes difficult. Likewise in the file treemap in Fig. 13, using colors to categorize too many types of files makes the information seem too overwhelming to even examine a subset. By reducing or more compactly binning the amount of information, the user may actually have a more efficient or accurate perceptual representation of the display.

5.1 Methods

Five subjects participated in this experiment. Three were female. All were either graduate or post-graduate students in psychology or computer science or trained university staff.

Fig. 9 shows the procedure for the subitizing task. We presented two different stimuli sequentially for 500 ms with a one second blank gray screen in between. The amount of variety in the two displays differed by one, and the order of smaller and larger displays was randomized. The probability of the first or second interval having more variety was even. The experiment had a two-interval forced choice (2IFC) design, so the subjects simply responded via the keyboard as to which interval had more variety. Due to the timed nature of the procedure, only the accuracy was recorded. Each subject performed 880 trials (4,400 total trials).

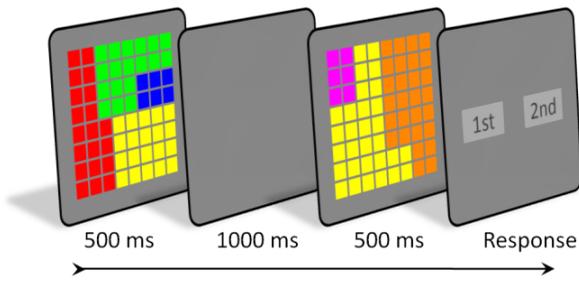


Fig. 9. The trial time course for the subitizing experiment. The first stimulus was presented for 500 ms followed by a 1000 millisecond gray screen and then the second stimulus for 500 ms. The subjects then responded in a two-interval-forced-choice which of the two intervals had more variety (e.g., more colors). Since the timing was dictated by the procedure, only accuracy was recorded rather than RTs.

5.2 Results

Fig. 10 shows the results trends (ANOVA results included below):

Accuracy and Variety: In contrast to the other experiments, none of the blocks could maintain a high accuracy as variety increased. All accuracies were nearly 100% for one variety but followed a marked, sigmoidal fall-off as variety increased. As additional detail was added, the subjects’ ability to understand the structure was severely impaired – $F(1,4) = 96, p < 0.005$.

Visual feature: Using 80% accuracy as a benchmark, the capacity for color variety across subjects was significantly higher at nearly double that of motion for grouped ($F[1,4] = 44, p < 0.0001$) and random layout ($F[1,4] = 18, p < 0.001$). This result is not surprising. Although the attentional capacity of motion with respect variety and layout has little research, color is known to have a high capacity [23], [24].

Layout: Again using 80% accuracy as a benchmark, the grouped layout had roughly 60% higher capacity for color ($F[1,5] = 44, p < 0.0001$) and motion ($F[1,3] = 18, p < 0.001$). A possible explanation is that visual working memory operates over or is modulated by grouped features. When the layout is complex, more cognitive resources are required thereby leaving less available for visual feature information. Irrespective of the psychological explanation, the data shows that if understanding the overall

structure of a visualization is important, a high-capacity visual feature and strong element clustering are critical.

Subject: Though the between-subject variability was significant, all subjects had the same trend of results (Fig. 10 error bars).

Edwards and Greenwood [25] performed a similar 2IFC experiment of subitizing motion presented via random dot stimuli (rather than moving texture patches). Their experiment, however, used the same number of elements for each variety on the display. Consequently, the quantity of any one feature type informed as to the number of variants in the stimulus. We therefore purposefully varied the quantity in this and the other experiments. Our subitizing motion results corroborated theirs, demonstrating that the capacity limits are based on the global or scene-wide information rather than the properties of a single element.

While our experiment focused on the accuracy of a quick glimpse, Watson et al [26] performed similar subitizing experiments examining how reaction time varied when unconstrained. Combined, our accuracy results and their response time results show an overall degradation of performance without grouping.

6 INFORMATION VISUALIZATION GUIDELINES

There are several particular results from our experiments that can inform visualization design:

6.1 Grouping greatly helps for some but not all tasks

When searching for an *unknown* oddball target—a target that may or may not, ostensibly, pop-out—the correlation between the spatial and visual feature axes is pivotal. High correlation speeds search; a cluttered randomly organized arrangement impairs search for the target. The degree to which the targets pop-out in Fig. 1 and Fig. 12 is modulated by the extent to which all the other colors are *grouped*.

Though the uses of information visualization can vary widely, we recommend prioritizing the tasks for which visualizations greatly outperform raw data (tables and databases). Extracting summary or contextual information, such as how much variety or how many categories of information are present is difficult with a spreadsheet, and doing so with a visualization relies on attention and visual memory, necessitating spatial grouping by visual feature. By incorporating visual feature into a visualization’s layout algorithm, observers could better judge overall qualitative characteristics of the display, such as whether the market had an overall gain or loss or how the success was distributed across industries (Fig. 11), or which

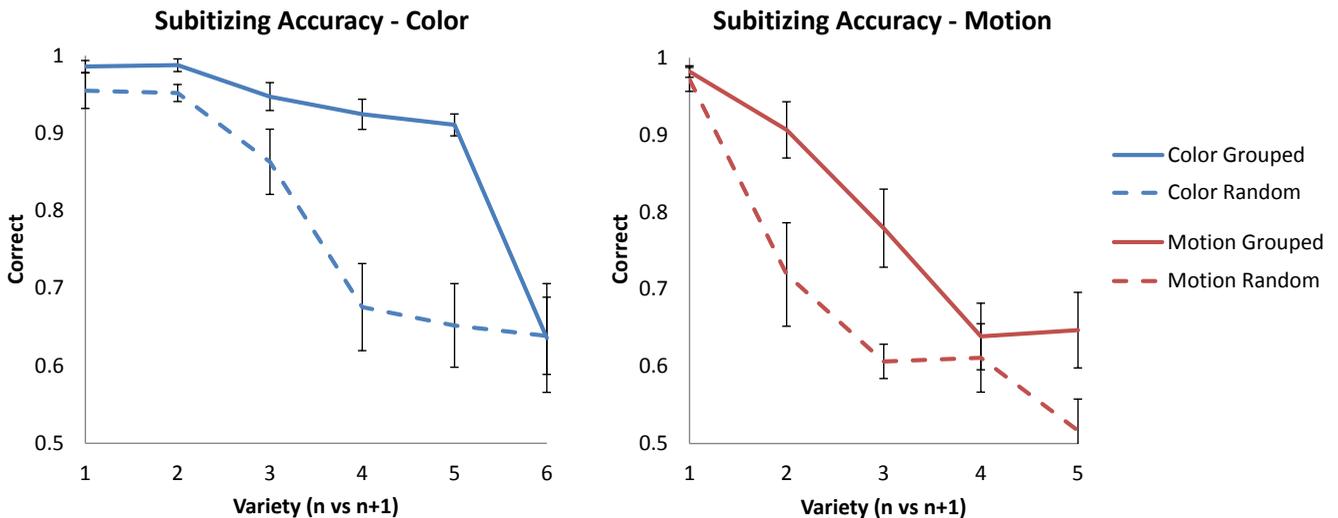


Fig. 10. The accuracy results of the subitizing experiment follow sigmoidal shape. The variety amount on the horizontal axis is the smaller of the two stimuli (1 means a comparison between 1 and 2 or 2 and 1). All layouts have nearly 100% accuracy for a comparison with one variety, and accuracy falls towards chance (50%) as variety increases. Grouping the features (color or motion) has a dramatic improvement on performance. Error bars show the inter-subject standard deviation.

industry had the most variance in performance.

This principle is a matter of prioritization, as optimizations for different visual features will likely conflict with each other. A color-focused layout would likely yield extreme aspect ratios which could hinder performance on size oriented tasks [27]. A visualization designer must therefore prioritize the visual feature that benefits the more important tasks.

6.2 If you cannot group, change the task

Knowing a target's appearance in advance dramatically improves accuracy and speed of target detection. Knowing to look for a red or rightward moving object drastically improves user performance, even within a highly variable display, suggesting that guided search is an important factor [8], [28]. The intuitive implication is that a legend, key, or any mechanism that helps a user know the visual appearance of a sought after data element greatly simplifies locating it.

6.3 When there are many categories: Less is more

A compromise must be made between the number of nominal categories and the perceptual complexity of a visualization. Design options are still available if a high number of nominal categories is needed. Though not always feasible, the spatial axes can be assigned to data dimensions that correlate with the visual feature's dimension [16]. As a result, the feature categories would become more spatially grouped and user performance would be less limited by the capacity. Because correlation is not always known, a perhaps more deterministic alternative is to limit the visualization to only show a couple of categories at a time (Fig. 13). The user would be required to interact to see different categories, but the limited information on the display could be analyzed within the limits of attention, enabling more efficient comprehension of the visualization.

6.4 Assigning a visual feature to a data dimension

We chose color and motion because we expected their performance and capacity to be fairly different (high for color and low for motion), and our result confirmed that hypothesis. Features should be chosen with care; features like color that have higher capacity should be reserved for data with high dynamic range or many categories.

6.5 Evaluation

Visual search for a known target is not always a sufficient test of visualization effectiveness. Though it may help assess how quickly certain kinds of information are conveyed in a visualization, it provides little or no information about the influence of capacity and other limits of attention. Evaluations of visualizations should make certain to test user performance for more attentionally demanding user goals. While we used oddball and subitizing tasks to examine attention capacity, other attentional limitations (e.g., the spatial and temporal resolution of attention [29], [30]) should be investigated depending on the users' goals. For example, if comparison between two datasets is the aim of a visualization, the impact of visual short term memory (e.g., change blindness, the inability to detect differences between spatially or temporally adjacent scenes [31], [32]) should be mitigated [33], [34]. This effect can be operationalized using visual short-term memory tasks including change detection [23], [24]. A limit on our ability to attend to multiple locations is another example of a potential hurdle for visualizations. A visualization that incorporates animations or moving objects should take into account the limits of human multiple object tracking [35]. In general, visualization designers need to test visualizations not only for simple pre-attentive tasks but also for tasks which are limited by attention.

7 FUTURE WORK

We showed the performance impact of grouping on a single, highly

discriminable categorical dimension. However, for a scalar, rather than categorical, dimension, the discriminability of the features can be an additional limiting factor. Further, data sets often represent multiple dimensions using more than one feature at a time (e.g., color and size). Examining and understanding how such multiplexing may limit capacity and search in the context of common visualization tasks could help the community understand how to more effectively visualize multiple dimensions.

8 CONCLUSION

We have three main conclusions: (1) Grouping is far more beneficial for oddball search compared with known-target search. (2) Accessing overall information (like heterogeneity or number of categories) is better for grouped displays. (3) For difficult tasks, aim to reduce variety in the entire view rather than optimizing small regions.

The implication is that the strict limits of attention have profound effects on the ability of observers to extract information from displays. Even performance with a visual feature like color, commonly thought of as 'pre-attentive', can be adversely effected by tasks or arrangements that put a heavy demand on attention and capacity. Together, there is an interaction between perceptual and cognitive limits and task demands. Accounting for these interactions can help design more efficient and effective visualizations.

ACKNOWLEDGMENTS

We would like to thank our experiment participants.

REFERENCES

- [1] A. Treisman, "A feature-integration theory of attention," *Cognitive Psychology*, vol. 12, no. 1, pp. 97-136, Jan. 1980.
- [2] J. Bertin, *Semiology of graphics: diagrams, networks, maps*. University of Wisconsin Press, 1983.
- [3] D. DiBiase, A. M. MacEachren, J. B. Krygier, and C. Reeves, "Animation and the role of map design in scientific visualization," *Cartography and Geographic Information Systems*, vol. 19, no. 4, pp. 201-214, 265-266, 1992.
- [4] L. Bartram, C. Ware, and T. Calvert, "Moticons: detection, distraction and task," *International Journal of Human-Computer Studies*, vol. 58, no. 5, pp. 515-545, May 2003.
- [5] D. E. Huber and C. G. Healey, "Visualizing Data with Motion," in *IEEE Visualization 2005 - (VIS'05)*, 2005, pp. 67-67.
- [6] W. S. Cleveland and R. McGill, "Graphical perception and graphical methods for analyzing scientific data," *Science (New York, N.Y.)*, vol. 229, no. 4716, pp. 828-33, Aug. 1985.
- [7] T. S. Horowitz and J. M. Wolfe, "Visual search has no memory," *Nature*, vol. 394, no. 6693, pp. 575-577, 1998.
- [8] J. M. Wolfe and T. S. Horowitz, "What attributes guide the deployment of visual attention and how do they do it?," *Nature reviews. Neuroscience*, vol. 5, no. 6, pp. 495-501, Jun. 2004.
- [9] J.-D. Fekete, J. J. V. Wijk, J. T. Stasko, and C. North, "The Value of Information Visualization," *Information Visualization*, vol. 4950, pp. 1-18, 2008.
- [10] E. L. Kaufman and M. W. Lord, "The discrimination of visual number," *The American journal of psychology*, vol. 62, no. 4, pp. 498-525, Oct. 1949.
- [11] J. Somervell, D. S. McCrickard, C. North, and M. Shukla, "An Evaluation of Information Visualization in Attention-Limited Environments," in *VISSYM 2002*, 2002, pp. 211-216.
- [12] J. Duncan and G. W. Humphreys, "Visual search and stimulus similarity," *Psychological Review*, vol. 96, no. 3, pp. 433-458, 1989.
- [13] A. Treisman, "Perceptual grouping and attention in visual search for features and for objects," *Journal of Experimental Psychology Human Perception and Performance*, vol. 8, no. 2, pp. 194-214, 1982.
- [14] S. Haroz and K.-L. Ma, "Natural Visualizations," in *Eurovis 2006*, 2006, pp. 43-50.
- [15] R. Rosenholtz, N. R. Twarog, N. Schinkel-Bielefeld, and M.

Wattenberg, "An intuitive model of perceptual grouping for HCI design," in *Proceedings of the 27th international conference on Human factors in computing systems - CHI '09*, 2009, p. 1331.

[16] A. Tatu, P. Bak, E. Bertini, D. Keim, and J. Schneidewind, "Visual quality metrics and human perception," in *Proceedings of the International Conference on Advanced Visual Interfaces - AVI '10*, 2010, p. 49.

[17] C. Ware and J. C. Beatty, "Using color dimensions to display data dimensions," *Human Factors*, vol. 30, no. 2, pp. 127-142, 1988.

[18] C. G. Healey, "Choosing effective colours for data visualization," in *Proceedings of Seventh Annual IEEE Visualization '96*, 1996, pp. 263-270.

[19] H. C. Nothdurft, "The role of features in preattentive vision: comparison of orientation, motion and color cues.," *Vision Research*, vol. 33, no. 14, pp. 1937-1958, 1993.

[20] L. Huang and H. Pashler, "A Boolean map theory of visual attention," *Psychological Review*, vol. 114, no. 3, pp. 599-631, 2007.

[21] D. Ariely, "Seeing Sets: Representation by Statistical Properties," *Psychological Science*, vol. 12, no. 2, pp. 157-162, Mar. 2001.

[22] S. C. Chong and A. Treisman, "Representation of statistical properties.," *Vision Research*, vol. 43, no. 4, pp. 393-404, 2003.

[23] G. A. Alvarez and P. Cavanagh, "The capacity of visual short-term memory is set both by visual information load and by number of objects," *Psychological Science*, vol. 15, no. 2, pp. 106-111, Feb. 2004.

[24] S. J. Luck and E. K. Vogel, "The capacity of visual working memory for features and conjunctions," *Nature*, vol. 390, no. 6657, pp. 279-81, Nov. 1997.

[25] J. A. Greenwood and M. Edwards, "The detection of multiple global directions: capacity limits with spatially segregated and transparent-

motion signals," *Journal of vision*, vol. 9, no. 1, pp. 40.1-15, Jan. 2009.

[26] D. G. Watson, E. A. Maylor, and L. A. M. Bruce, "The efficiency of feature-based subitization and counting.," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 31, no. 6, pp. 1449-1462, 2005.

[27] N. Kong, J. Heer, and M. Agrawala, "Perceptual guidelines for creating rectangular treemaps.," *IEEE transactions on visualization and computer graphics*, vol. 16, no. 6, pp. 990-8, Jan. 2010.

[28] J. M. Wolfe, "Guided search 2.0: a revised model of visual search," *Psychonomic Bulletin & Review*, vol. 1, no. 2, pp. 202-238, 1994.

[29] J. Intriligator and J. P. Cavanagh, "The Spatial Resolution of Visual Attention," *Cognitive Psychology*, vol. 43, no. 3, pp. 171-216, 2001.

[30] A. O. Holcombe, "Seeing slow and seeing fast: two limits on perception.," *Trends in cognitive sciences*, vol. 13, no. 5, pp. 216-21, May 2009.

[31] R. A. Rensink, J. K. O'Regan, and J. J. Clark, "To See or not to See: The Need for Attention to Perceive Changes in Scenes," *Psychological Science*, vol. 8, no. 5, pp. 368-373, Sep. 1997.

[32] D. J. Simons and R. A. Rensink, "Change blindness: past, present, and future," *Trends in cognitive sciences*, vol. 9, no. 1, pp. 16-20, Jan. 2005.

[33] L. Nowell, E. Hetzler, and T. Tanasse, *Change blindness in information visualization: a case study*. IEEE, 2001, pp. 15-22.

[34] S. Haroz and K. Heitmann, "Seeing the difference between cosmological simulations," *IEEE computer graphics and applications*, vol. 28, no. 5, pp. 37-45, Sep. 2008.

[35] Z. W. Pylyshyn and R. W. Storm, "Tracking multiple independent targets: Evidence for a parallel tracking mechanism," *Spatial Vision*, vol. 3, no. 3, pp. 179-197, 1988.

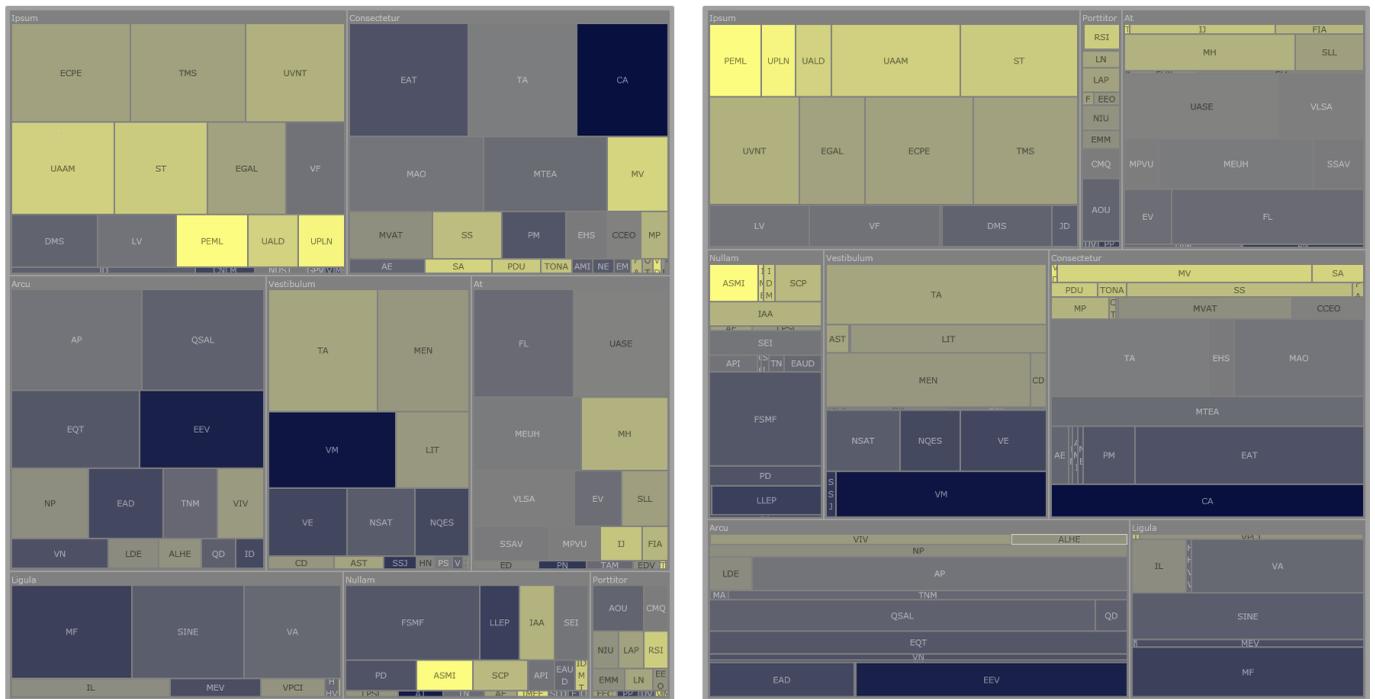


Fig. 11. In visualizing a stockmarket using a sorted squarified treemap (left), the impact of prioritizing aspect ratio is a scattering of the colored rectangles. By prioritizing color and size (right), better visual grouping occurs, and more trends can be found by users. In these two stock market visualizations, company capital corresponds to size, color is a function of intraday percent change, and the stocks are segmented by industry. This is a similar approach used by smartmoney.com. Can you tell which industry performed best? Or worst? (hint: the two visualization use the same data)



Fig. 12. The left image with news from multiple countries is technically presenting more information to the user. Yet the complexity of the layout severely degrades basic information (which news category is prevalent?). On the right the grouping that occurs when only one country is selected simplifies any analysis of the proportions between news categories. Though they both have the same number of categories and (due to screen real estate) are only capable of utilizing the same fixed area for stories, the right image appears more streamlined and useful.

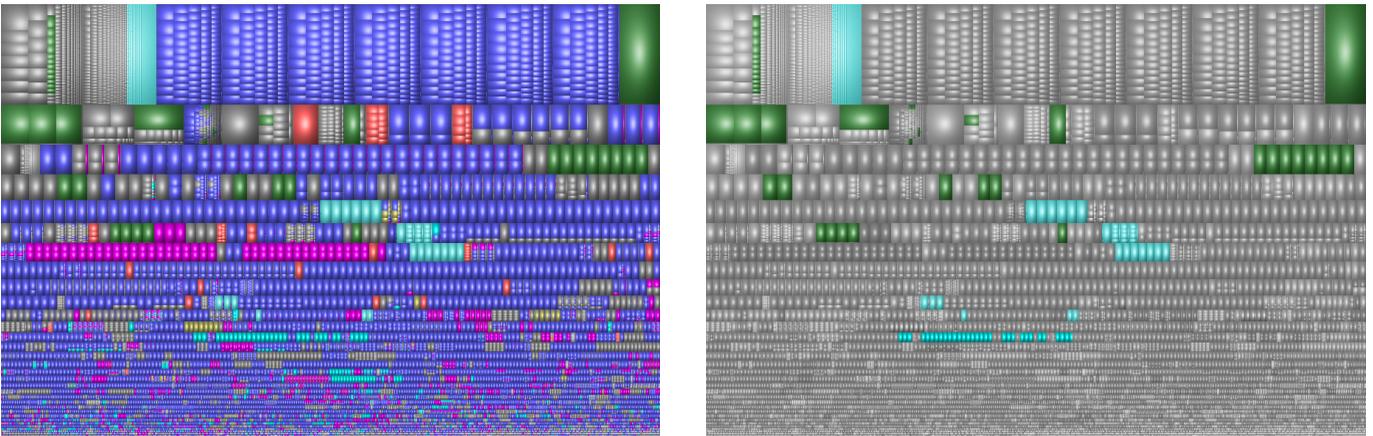


Fig. 13. (left) This image of WinDirStat shows a complex arrangement of many types of files. Trying to grasp the number of categories present or their relative organization relative to each other is difficult unless attention is focused on a single variant of a feature (attend to green) or the filtration is performed for the user (right).