Identifying Best Practices for Visualizing Photo Statistics and Galleries using Treemaps

John A. Guerra-Gomez
Yahoo Labs
john.guerra@gmail.com

Cati Boulanger
Yahoo Labs
catinettev@gmail.com

Sanjay Kairam
Stanford University
skairam@cs.stanford.edu

David A. Shamma
Yahoo Labs
aymans@acm.org

ABSTRACT

Online photo collections are often presented along with numeric data, such as views, likes, or comments. Treemaps are ideal for visualizing such collections, as they present numeric values using rectangular blocks, within which photos can be presented. Despite abundant research showing that photo treemaps and similar space-filling approaches are useful and appealing for users, understanding of the ideal parameters for constructing these representations remains limited. To address this, we contribute a series of experiments targeted at identifying design parameters for building photo treemaps. Our first study explores the number of photos presented at each level in the treemap, finding that using fewer photos makes visualizations more searchable and preferable to users. A second study examines the tradeoff between sizing photos according to the numeric values and presenting them in a more familiar layout with fixed aspect ratios. These results inform a third study, in which a prototype was used to probe the preferences of active Flickr users for using treemaps for navigating photo collections.

CCS Concepts

•**Human-centered computing** → **Treemaps; Information visualization;** *Field studies;*

Keywords

Information Visualization; Photo Visulizations; Photo treemaps

1. INTRODUCTION

Since their introduction in 1992 [12], treemaps have been widely accepted as a powerful and useful technique for visualizing hierarchical and non-hierarchical datasets. Photo treemaps, as first introduced in PhotoMesa [1, 9], leverage the hierarchy-preserving features of treemaps to represent the folder or category structure of photo collections. Since then, many systems have adopted similar space-filling layouts for representing collections [2–8, 10, 11, 13, 14].

AVI '16, June 07 - 10, 2016, Bari, Italy

ISBN 978-1-4503-4131-8/16/06. . . $15.00

DOI: http://dx.doi.org/10.1145/2909132.2909280

Figure 1: Screenshot of experiment 1, where turkers were asked to select which widget they prefer between two treemaps with a different number of photos responses.

Photos in online collections are often presented along with associated numeric values, frequently representing markers of social activity (e.g. views, likes, favorites, or comments). These associated data can communicate valuable signals to users exploring collections about photos which may be of particular interest or high quality. In a blog post [1], Tashian first introduced the idea of laying out a photo treemap so that the size represented the relative importance of a particular photo, according to a chosen metric. In this paper, we refer to this technique as *proportional photo treemaps*.

In constructing his proportional photo treemap, Tashian expressed difficulties with the layout algorithm, citing issues with aspect ratios and cropping. Automating layout for proportional photo treemaps requires careful balancing of representation and aesthetic concerns. More faithful representation of data values may make it increasingly difficult to render photos without cropping or transformations. Choosing the correct design parameters is both difficult and important for providing a high-quality user experience.

In this paper, we address these challenges through three targeted user interface studies aimed at identifying ideal parameters for designing photo treemaps. In the first study, we explored how user preferences varied with the number of photos represented in the visualization, finding that users tended generally to prefer fewer photos. In the second study, we compared proportional photo treemaps to a modified scrolling layout which rendered all photos with a fixed aspect ratio and constant size. Despite our hypothesis that the

[1] http://tashian.com/carl/archives/2004/06/data_visualizat.php

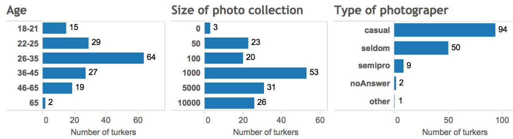

(a) Experiment 1.

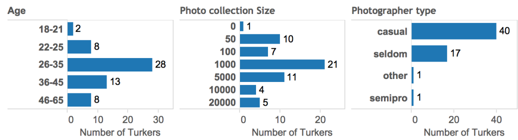

(b) Experiment 2.

Figure 2: Demographic information for participants in the first two experiments.

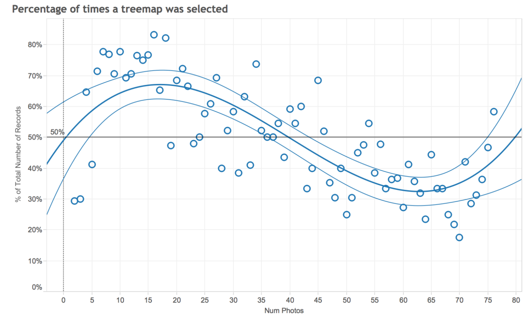

Figure 3: Percentage of time a treemap was selected according to the number of photos it displayed in experiment 1. The results show that the ideal number of photos for a treemap of about 390×440px is between 5 and 25. The cubic regression with its confidence intervals is shown, $y = 7.21{*}10^{-6}x^3 - 0.000865x^2 + 0.0232653x + 0.490336$, $p < 0.0001$.

scrolling layout would be more difficult to navigate, we found that users could find photos more quickly and accurately using this layout. Finally, we ran a series of semi-structured interviews with users of a popular photo service (Flickr); in this study, they used a prototype system to browse their personal repositories, providing insights into their requirements and preferences around layout, sorting, and labeling.

The rest of this paper is organized as follows: we first present a brief review of the literature on the use of treemaps and other techniques for representing photo collections. We then describe in detail each of the three experiments. We conclude with some discussion about the findings and how they can be applied to both the design and future study of photo treemap visualizations.

2. EXPERIMENT 1: PHOTO COUNT

In this first study, we evaluated how user preferences varied with the number of photos displayed in a proportional photo treemap visualization.

Procedure. We recruited participants using Amazon Mechanical Turk and the psiTurk framework. HITs were posted over the course of three days, accepting only workers in the U.S. with more than 95% positive feedback. We received 154 responses (67 women), with demographics as presented in Figure 2(a). Participants received $0.75 per HIT; each HIT took a median of 68 seconds to complete, resulting in an approximate wage of $39 per hour.

Each HIT consisted of five trials. In each trial, participants were presented with an interface that showed two treemaps side-by-side. In each trial, the two treemaps were presented with equal dimensions (390×440px), but with different numbers of photos (selected at random from the range [2,76]). The photos represented in the treemaps were drawn from Flickr's most popular public photos of the day. In each trial, the participant first completed a quality control task, in which they were asked to find a specified photo in each treemap. They then were asked to choose which of the two treemaps they preferred. After excluding three incomplete responses, we were left with 755 completed trials from 151 participants, which we discuss below.

Results. The main outcome variable of interest was users' preference as a function of the number of photos in the layout. As illustrated in Figure 4, we found that participants preferred the treemap with fewer photos roughly 64% of the time, with a binomial test indicating that this result was unlikely to happen by chance, with $p < 0.001$ (two-tailed). A more detailed analysis showed that participants generally preferred a layout with 5 to 25 photos given the size of the treemap presented (390×440 pixels).

Though the photo-finding task was designed as a quality control measure, analyzing the results revealed some interesting findings. In each treemap, the ranking of a photo was doubly-encoded using the position and size, with highly-ranked photos displayed larger and in the top-left corner. To explore the factors which influenced finding times, we built a linear mixed-effects model with photo-ranking, interface side (left or right), and number of photos in the layout as fixed effects, with a random intercept for participant (representing that participants have different aptitudes for the task). Because our response variable, finding time, was log-normally distributed, we used a log transformation.

To assess how each of the three main variables impacted our model, we compared this full model to three reduced models in which one of these variables had been removed, using likelihood-ratio tests. We observed a significant effect of photo ranking on finding time ($\chi^2(1) = 42.528$, $p < 0.001$), with each increase in rank increasing finding time by roughly 1.1%. As expected, we also observed a significant effect of photo count ($\chi^2(1) = 91.762$, $p < 0.001$); each additional photo in the layout increased finding time by 0.9%.

Additionally, we observed a significant effect of interface side (left vs. right) on finding time ($\chi^2(1) = 27.971$, $p < 0.001$), with photos taking roughly 20% longer to find in the treemap on the right side. This phenomenon may be explained by an animation which expanded photos on mouse hover; as photos expanded towards the bottom right, photos in the right treemap could potentially expand beyond the rightmost bound of the screen.

3. EXPERIMENT 2: LAYOUT

In the second study, we explored the tradeoffs between proportional photo treemaps (where photos are sized according to associated numeric values) and an alternate scrollable version where photo sizes were fixed.

Procedure. We again recruited participants through Mechanical Turk, using the same screening criteria. We received 63 responses (27 women), with demographics as presented in

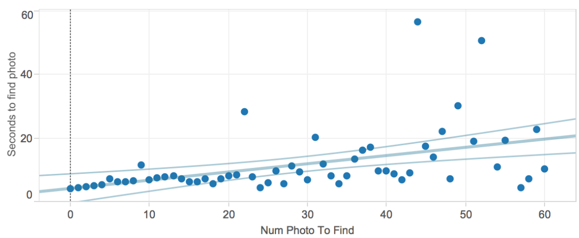

Figure 4: Median time to find a photo by ranking in the experiment 1; the photo in position 0 is the largest and positioned in the top-left, making it easier to find. The linear regression with its confidence intervals is shown, $y = 4.12604x + 0.257939$, $p = 0.0003727$.

Figure 2(b). Participants received $0.75 per HIT; each HIT took a median of 165 seconds to complete, resulting in an approximate hourly wage of roughly $16.

Each HIT consisted of six trials, three for each layout condition. In each trial, participants were presented with an interface showing one of the layouts, visualizing 48 photos in a space of 536×588px. In the proportional layout, photos were sized according to importance, compressing all photos into the available space. In the scrolling layout, photos extending beyond the available space were accessible by scrolling.

Each trial consisted of two finding tasks; in each task, the participant were presented with a single photo and informed that their task would be to find this photo in the subsequent screen. When they had familiarized themselves with the photo, they clicked a button which displayed the layout; task time was measured from this button click until the time they found and selected the photo in the layout. Participants then performed a second finding task with this same layout. After completing both finding tasks, they were asked to select the photo from the layouts the photo that was more popular. After completing all trials, participants were asked to choose which layout they found more useful for the given task.

Results. After excluding 4 incomplete results, we were left with responses from 59 participants. To explore the factors which influenced finding times, we build another linear mixed-effects model with interface condition (scrolling vs proportional), photo position, and task order (first or second) as fixed effects, and with a random intercept for participant. Again, the response variable, time, was log-transformed. As in the previous study, we assessed the impact of each of the main variables by comparing the full model to reduced models in which one of these variables had been removed.

We observed that finding time varied significantly with interface condition ($\chi^2(1) = 33.734$, $p < 0.001$), with participants using the proportional view taking roughly 42% longer. We again observed a significant effect of photo position ($\chi^2(1) = 126.06$, $p < 0.001$); each increase in position increased finding time by roughly 2.6%. Finally, as expected, we observed that participants performed the second finding task in each trial significantly faster than the first ($\chi^2(1) = 16.069$, $p < 0.001$), doing so approximately 21% faster. Participants made a total of 21 errors in the scrolling condition and 47 in the proportional condition. In terms of overall preferences, 48/59 (81.3%) participants favored the scrollable layout over the proportional layout.

In response to the question about which photo was more

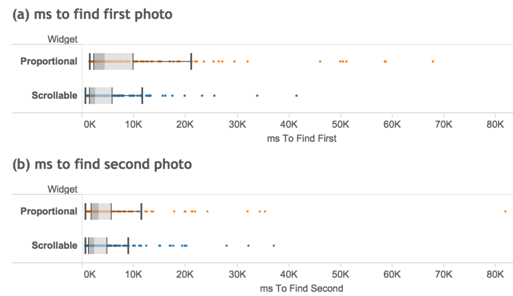

Figure 5: Experiment 2 results. Turkers were faster with the scrollable widget compared to the proportional treemap, and faster in finding the second photo for each layout. Although there are 2 outliers not displayed for clarity, they were considered in the box plot computations (e.g. One turker spent 409s finding the first photo in the proportional treemap).

"popular", participants chose the correct photo 153/177 times (86.44%) when using the scrolling layout and 138/177 (77.97%) when using the proportional layout. To explore the factors which influenced success in choosing the more popular photo, we used a logistic model with two fixed-effects (interface and difference in ranking between the two photos) and one random intercept (participant). Interface had a significant effect on choosing correctly ($z = -2.323$, $p = 0.202$); the odds of success in the scrolling view are 109% higher than in the treemap view. We also observed a significant effect of the difference in photo positions ($z = 4.045$, $p < 0.001$); increasing the difference in photo positions by one increased the odds of choosing correctly increased by 7.7%. Judging by the feedback provided in the post-task questionnaire, only a handful of participants guessed when responding, most selected photos close to the top, but some preferred photos on the top right corner or even in the middle.

4. EXPERIMENT 3: USER STUDY

Our first two experiments explored specific parameters for designing photo treemaps. In the third study, our aim was to elicit users' more general requirements and preferences for photo treemaps. We conducted a user study with active users of a photo sharing service in order to understand how photo treemaps can aid photographers in gaining insight into social engagement around their own photos.

Procedure, We recruited seven participants through an advertisement posted on Craigslist; all participants completed a survey, hosted on SurveyMonkey, designed to elicit qualitative and quantitative feedback on a static treemap visualization. We received a total of 7 responses.

Participants were presented with a series of six proportional photo treemap visualizations, each constructed with 40 of the user's most popular photos. The treemaps presented varied along two dimensions: 1) the choice of numeric variable for sizing and sorting (views, favorites, or comments), and 2) the presence or absence of labels inside the boxes. For each image, we elicited feedback in the form of open-ended questions and two 5-point Likert scale questions measuring pleasingness and usefulness. After viewing all of the configurations, a final question asked participants to rank all six. All open-ended responses were hand-coded

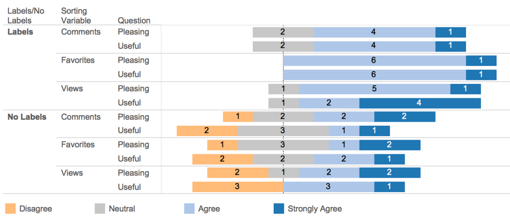

Figure 6: Experiment 3 quantitative responses in the form of 5-Likert scales. Participants show strong preference for labeled configurations of the photo treemaps for all sorting conditions and both questions.

to identify common themes across respondents, comments comparing the different configurations, and potential areas for improvement.

Results. All (7/7) of the participants preferred the presence of labels, as they helped make sense of the photo representation; all found the variations with labels to be intuitive and engaging. They preferred viewing these labels (showing the number of views, faves, or comments) directly on the photos itself, finding these views not only more pleasing but also more useful. User preferences as expressed in the Likert-scale questions are summarized in Figure 6. In terms of ranking, participants selected overall in this order: 1) Most viewed with labels, 2) most favoured with labels, 3) most viewed without labels, 4) most commented with labels, 5) most favoured without labels and finally 6) most commented without labels.

Participants proposed that if the image size relative to the number of views was more exaggerated, it would be more intuitive. One participant expressed that it gets too "crowded" when the photos are too small. Participants proposed to "hover" on the photo to see the label, instead of it being permanently on the photo. Participants appreciated the ordering presented in the treemaps, left/right or top bottom, which let them immediately visualize their most important photo,

The feedback from the participants depends highly on the activity on their photos. For instance, P1 explains: *"Given that I do not have lots of comments, seeing this visualization is not very useful to me personally"*, referring to the comments based treemap. Participants also explained that they would like to be able to access the comment itself and active links to the photos themselves.

5. DISCUSSION AND CONCLUSIONS

In this paper we evaluated treemaps as a technique for analyzing personal photo statistics and galleries. The results of experiment 1 indicate that although treemaps are capable of showing many photos in a reduced space, cropping and ratio restrictions have a big effect on users' preference. Albeit, users preferred generally smaller treemaps, there is a lower limit on the number of photos users want to see. For example, the sweet spot of photos in a treemap of 390×440px is about 10 to 25 photos.

Experiment 2 revealed that despite the hurdles of scrolling, a scrollable widget can be more appealing than a proportional photo treemaps. This suggests that photo cropping needs to be carefully addressed for proportional photo treemaps to work in real world scenarios. Despite this, when presented with photo treemaps of their own photos, users found them useful and appealing (as suggested by experiment 3). Further experiments are required to validate if photo ownership plays a role in user's preference of photo treemap layouts.

Finally, experiment 3 suggested that labels are important for helping users understand photo treemaps, even if is only while they get used to the interface. The proportional photo treemaps were more appealing for users with more photo activity. Order was very important for users and again photo cropping was a concern. For future work we want to investigate techniques for reducing photo cropping as well as evaluate the use of treemaps for comparing sets of photos, and for navigating hierarchical categorizations of photos.

6. REFERENCES

[1] B. Bederson. PhotoMesa: a zoomable image browser using quantum treemaps and bubblemaps. *UIST '01*.

[2] P. Brivio, M. Tarini, and P. Cignoni. Browsing large image datasets through Voronoi diagrams. *IEEE trans. on vis. and computer graphics*, 16(6):1261–70, 2010.

[3] P. Brivio, M. Tarini, F. Ponchio, P. Cignoni, and R. Scopigno. PileBars: Scalable Dynamic Thumbnail Bars. *VAST*, 2012.

[4] E. Cambria and A. Hussain. Sentic Album: Content-, Concept-, and Context-Based Online Personal Photo Management System. *Cognitive Computation*, 4(4):477–496, 2012.

[5] A. Girgensohn, J. Adcock, M. Cooper, J. Foote, and L. Wilcox. Simplifying the management of large photo collections. *Human-Computer Interaction INTERACT*, 3:196–203, 2003.

[6] O. Hilliges, D. Baur, and A. Butz. Photohelix: Browsing, Sorting and Sharing Digital Photo Collections. In *TABLETOP'07*, pages 87–94. IEEE.

[7] D. F. Huynh, S. M. Drucker, P. Baudisch, and C. Wong. Time quilt. In *CHI '05*, page 1937, New York, New York, USA. ACM Press.

[8] K. Karlsson, W. Jiang, and D.-Q. Zhang. Mobile Photo Album Management with Multiscale Timeline. In *MM '14*, pages 1061–1064. ACM Press.

[9] A. Khella and B. B. Bederson. Pocket PhotoMesa. In *MUM '04*, pages 19–24, New York, New York, USA, 2004. ACM Press.

[10] P. O. Kristensson, O. Arnell, A. Björk, N. Dahlbäck, J. Pennerup, E. Prytz, J. Wikman, and N. Åström. InfoTouch: an explorative multi-touch visualization interface for tagged photo collections. In *NordiCHI '08*, page 491, New York, New York, USA. ACM Press.

[11] J. Moehrmann and S. Bernstein. *Improving the usability of hierarchical representations for interactively labeling large image data sets.* Springer Berlin Heidelberg, 2011.

[12] B. Shneiderman. Tree visualization with tree-maps: 2-d space-filling approach. *ACM Transactions on Graphics*, 11(1):92–99, jan 1992.

[13] C. Wang, J. P. Reese, H. Zhang, J. Tao, and R. J. Nemiroff. iMap: a stable layout for navigating large image collections with embedded search. In *IS&T/SPIE Electronic Imaging*, page 86540K, 2013.

[14] J. Xiao and T. Zhang. Face bubble. In *AVI '08*, page 343, 2008.