

Comparison of Tag Cloud Layouts: Task-related Performance and Visual Exploration

Steffen Lohmann, Jürgen Ziegler, and Lena Tetzlaff

University of Duisburg-Essen,
Lotharstrasse 65, 47057 Duisburg, Germany,
{steffen.lohmann, juergen.ziegler, lena.tetzlaff}@uni-due.de

Abstract. Tag clouds have become a popular visualization and navigation interface on the Web. Despite their popularity, little is known about tag cloud perception and performance with respect to different user goals. This paper presents results from a comparative study of several tag cloud layouts. The results show differences in task performance, leading to the conclusion that interface designers should carefully select the appropriate tag cloud layout according to the expected user goals. Furthermore, the analysis of eye tracking data provides insights into the visual exploration strategies of tag cloud users.

Keywords: Tag Clouds, Tagging, Performance, Perception, Eye Tracking, Visualization, User Goals, Navigation, Evaluation, Folksonomy.

1 Introduction

With the advent and great success of a new generation of community-oriented websites in domains such as Media Sharing (e.g., Flickr, YouTube) or Social Bookmarking and Citation (e.g., Delicious, Connotea), a new way of metadata creation has emerged, commonly known as *tagging*. Tagging-based systems enable users to add *tags* – freely chosen keywords – to Web resources to organize these resources for themselves and/or others. Visual browsing in the tag collections is realized in different ways; often, websites offer an interface element known as *tag cloud*. Usually, a tag cloud presents a certain number of most often used tags in a defined area of the user interface. A tag's popularity is expressed by its font size (relative to the other tags) and is therefore easily recognized. Sometimes, further visual properties, such as the font color, intensity, or weight, are manipulated (for an overview see [2]). Next to their visualization function, tag clouds are also navigation interfaces as the tags are usually hyperlinks leading to a collection of items they are associated with.

However, tag clouds are not only used to display tag sets but are also increasingly applied in other contexts and for various data sets, for instance, in the areas of information visualization or text summarization (cp. [11]). Furthermore, several layout variations emerged on the basic design principles of tag clouds. Most popular is the 'classic' rectangular tag arrangement with alphabetical sorting in a sequential line-by-line layout. Fig. 1a shows an example of this well-known type of tag cloud

from the photo sharing website Flickr [12]. One reason for the popularity of this tag cloud layout might be its easy implementation. However, a sequential arrangement of tags may not provide optimal support for all purposes a user consults a tag cloud for. A wide range of alternative layouts, such as circular or clustered tag clouds (see Fig. 1b), have been proposed.

Although little is known about the benefits of tag clouds and the situations in which website visitors use them, some typical user goals have been identified – from getting an overview or general impression of a website’s contents to searching for specific topics or tags (cp. [5, 9, 10]). It has also been conjectured that tag clouds serve a social purpose by conveying a sense of activity in a Web community [7]. The ability of different tag cloud layouts to support certain user goals and information needs, however, is still a largely open research issue.

This paper takes a closer look at tag cloud performance and perception by investigating different tag cloud layouts in their ability to support users in typical information seeking tasks. In Sec. 2, we give a short overview of related work and summarize recent research findings. We then describe the experimental setup and the systematic generation of tag clouds, user tasks, and scenarios in Sec. 3, followed by a description of the experimental procedure in Sec. 4. Sec. 5 and 6 report and discuss results of task-layout-performance and perceptual patterns that could be analyzed from the eye tracking data. Based on these results, we draw some general conclusions in Sec. 7.

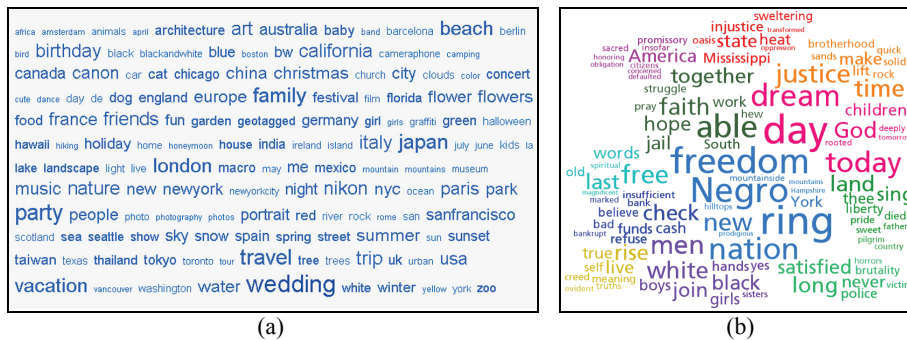


Fig. 1. Tag cloud examples: (a) sequential layout with alphabetical sorting, showing the “all time most popular tags” of the photo sharing website Flickr [12], (b) statistically clustered tag cloud, presenting keywords of the famous speech “I Have a Dream” by M. L. King Jr. [3].

2 Related Work

Most existing studies evaluate tag clouds at a general level by comparing them with other types of user interfaces. For instance, Halvey and Keane compared tag clouds with unweighted, horizontal and vertical lists by asking participants to find and select specific tags in both interface types [5]. The results indicate that unweighted lists perform better than tag clouds and that alphabetical sorting further accelerates search. In addition, the tags’ font sizes had a strong effect on search speed and tags in the

upper left corner of the cloud were found most quickly. Based on the times the study participants needed to find tags in specific target regions of the cloud, the authors also concluded that tag clouds are rather scanned than read. Kuo et al. got similar performance results by comparing tag clouds and lists as alternative presentation forms for search result summarization in the biomedical domain [8]. Though the study participants solved the descriptive tasks more quickly with the lists they attributed the tag clouds a higher level of satisfaction.

Sinclair and Cardew-Hall [10] conducted an experiment in which the subjects could either use a tag cloud or a traditional search interface to answer questions. They found that the subjects preferred the search interface for specific information retrieval tasks whereas the tag cloud was preferred for more open-ended questions. They concluded that tag clouds are not a replacement but a valuable extension to other types of user interfaces.

Other work takes a closer look at the visual features of tag clouds. Bateman et al. varied eight tag cloud properties and measured their effects on tag selection [2]. Their results show that some properties, such as font size and font weight, have stronger effects on tag selection than others, such as the intensity of a tag's color or its number of characters. Furthermore, the authors observed a central tendency: Tags in the middle of the cloud were selected more often than tags in the top and bottom areas.

Rivadeneira et al. conducted two experiments [9]: The first examined the effects of font size, location, and proximity-to-the-largest-tag. The user task was to recall terms that were previously presented in tag clouds. In accordance with [5] and [2], a strong effect of font size was observed, while proximity-to-the-largest-tag had no significant impact. Furthermore, tags in the upper left quadrant were recalled more often, but the authors attributed this to the sparseness of the tag clouds that were used in the study. In the second experiment, Rivadeneira et al. investigated the effects of font size and tag cloud layout on impression formation and recognition. Again, font size had a strong effect on the results. The tag cloud layouts had no impact on recognition but affected the accuracy of impression formation. In accordance with [5] and [8] it showed that an unweighted, one-dimensional list performs slightly better than a tag cloud.

2.1 Discussion of Related Work

While the presented studies provide valuable findings on tag cloud performance and are first steps towards a better understanding of tag cloud perception, they mainly focused on a sequential layout without considering further ways to arrange tags. The only study that evaluates different tag cloud layouts [9] is very limited in its conclusions as these layouts were highly diverse leading to many interdependencies in the visual features. Moreover, layout performance was only examined regarding recognition and impression formation in this study but not with respect to more typical situations of tag cloud interaction.

In many of the studies, the experimental material was comparatively unusual and artificial. For instance, in some cases the tag clouds consisted of a very limited number of tags (e.g., only 10 tags [5] or 13 tags [9] in total) or few font size variations (e.g., only 3 different font sizes [5]). In other cases, the used tag corpus was rather untypical for Web contexts [2, 9].

Most important, none of the studies addressed the visual exploration of tag clouds directly. All conclusions regarding tag cloud perception have been indirectly inferred from the tags that were either selected or recalled by the study participants. None of the studies recorded any gaze data or considered visual exploration strategies.

3 Experimental Setup

In contrast to related work, we pursued the following objectives with our study:

- Defining user tasks that simulate situations of nearly realistic tag cloud interaction.
- Presenting tag clouds with a common number of tags and font size variations.
- Generating tag cloud layouts that differ only in tag arrangement but no other visual properties.
- Using eye tracking to measure the actual attention areas and perception patterns of tag clouds.

3.1 User Tasks

From the range of purposes tag clouds might be used for, we selected three search tasks that are very common according to the literature [2, 5, 8, 9, 10]:

1. Finding a specific tag
2. Finding the most popular tags
3. Finding tags that belong to a certain topic

By providing a scenario along with the tasks, a somewhat realistic interaction situation was simulated. For example, one text for the first task was: “You always wanted to visit Nizza. At the moment, you have not enough money for travelling and can only dream about it... Please click on the term ‘Nizza’ in the following tag cloud.” Thus, the participants were asked to search for a target tag and click on it in this task. In the second and third task, the participants had to search and select three tags of choice that were either among the most popular ones (task 2) or related to a given topic (task 3). All study participants reported that they had understood the tasks; none needed further explanation or complained in the post-test questionnaire.

3.2 Tag Cloud Layouts

In order to reduce the countless variations of possible tag arrangements to a manageable and reasonably representative number, we made the following assumption: Many user interfaces are based on a two-dimensional grid layout that conceptually divides the interface into a number of rectangles by horizontal and vertical lines [1]. Thus, a typical requirement for a balanced integration of tag clouds into user interfaces is that they (1) are displayed in a rectangular area and (2) fill this area with tags as completely as possible. Based on these requirements, three main ordering principles can be distinguished:

1. *Sequential layout*, with either a horizontal or vertical arrangement of tags, sorted alphabetically or by some other criteria (e.g., popularity, chronology, etc.)
2. *Circular layout*, with the most popular tags in the center and tags with decreasing popularities towards the borders (or vice versa)
3. *Clustered layout*, in which the distance between tags follows a certain clustering criteria (e.g., semantic relatedness) and related tags are positioned in close proximity [3, 6]

We generated one prototypical tag cloud for each of these three layout strategies. In addition, we created a fourth tag cloud with no variation in the tags' font sizes as reference layout (see Fig. 2).

3.3 Tag Corpora

Since we aimed to compare the different tag cloud layouts independently of interpersonal differences, we decided to present all four layouts in a series to every study participant. We used tag corpora from four different domains to prevent learning effects that could likely result from presenting the same tags in each layout.

Tag corpora from "real" Web contexts were not well suited for our controlled experiment as they usually contain terms with heterogeneous connotations and little semantic consistency. We aimed at minimizing any bias caused by personally affecting or political terms (such as "terrorist"). Furthermore, we wanted to present tags that are well-known to the study participants to avoid effects that result from different interests and educational backgrounds. Additional requirements for the tag corpora included a reasonable way of categorizing the tags for the clustered tag cloud and the usage of terms that are common in Germany, since the study participants were all German-speaking. For these reasons (and due to a lack of alternatives) we decided to develop our own tag corpora for the study that consisted of neutral terms from common knowledge areas (France, sports, furniture, animals) and could be used to create somewhat realistic interaction scenarios. Each corpus consisted of 100 tags with varying popularity values.

3.4 Generation of the Tag Clouds

Since our goal was to evaluate general layout types and not specific algorithms, we decided to generate the tag cloud layouts by our own instead of using available algorithms. That way, we could apply the same design guidelines for all tag clouds and strictly control tag distribution. Nevertheless, the tag clouds were designed to largely follow typical algorithms and were not perceived as artificial or unusual by the study participants.

We generated the tag clouds for all four layouts as follows: We used an equally sized rectangle with an aspect ratio of 3:2 and filled it with the 100 tags of one of the four corpora. We mapped the tags' popularity values on six discrete font sizes, resulting in one tag of 30 pt up to 27 tags of 15 pt for each tag cloud. Every quadrant of the tag cloud got the same number of tags of each font size to avoid biases caused by an unbalanced presentation. In accordance with the ordering principle of the

circular layout, the 30 pt tag was placed in the middle of the cloud. In the sequential and clustered layout, the 30 pt tag was placed in another quadrant for each of the four corpora. We also considered the different user tasks in the tags' distribution: We varied the size and quadrant position of the tags the participants were asked for in the first task. Likewise, we distributed the thematic clusters that were of interest in the third task among all four quadrants of the tag clouds. We kept all other visual features, such as font styles, weights, colors, or intensities, constant in order to avoid interdependencies as reported in [2].

Fig. 2 shows the resulting four layouts that were generated for the corpus 'France'. The colored lines, circles, and arrows indicate quadrant separations and ordering principles of the layouts (not visible for the study participants).

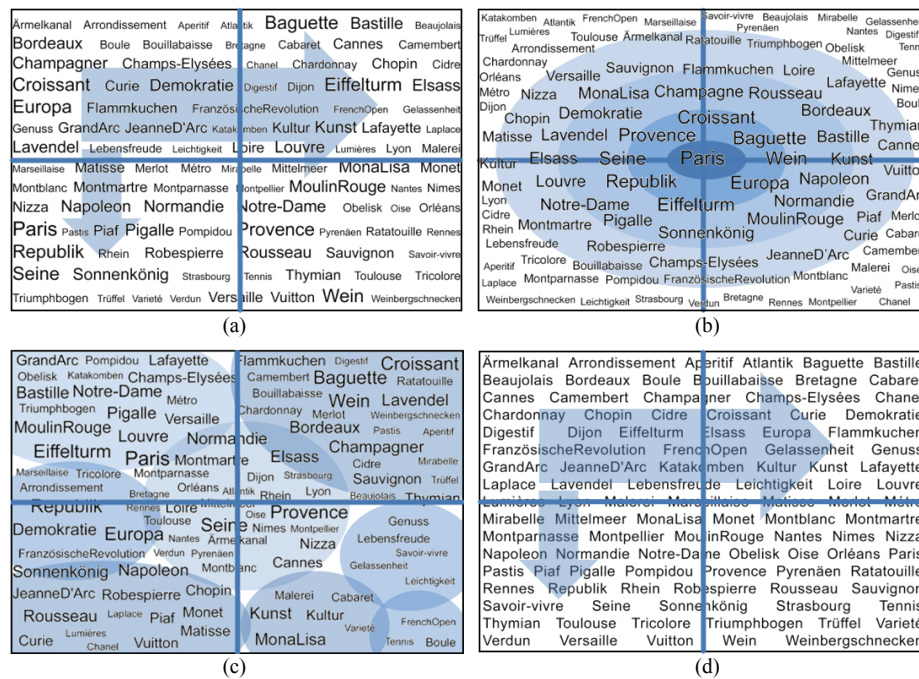


Fig. 2. Tag cloud layouts for the corpus 'France': (a) sequential (alphabetical sorting), (b) circular (decreasing popularity), (c) clustered (thematic clusters), (d) reference (sequential, alphabetical sorting, no weighting of tags).

4 Experimental Procedure

36 participants, mainly students, with an average age of 26 (min 17, max 57) took part in the experiment. All had normal or corrected-to-normal vision. The general familiarity with tag clouds was given with a median of 4 on a scale of 1 to 10. Five participants could not remember having seen any tag clouds before.

In order to create a nearly identical understanding of tagging and tag clouds, the experiment started with a short introduction into these topics, consisting of an oral explanation accompanied by a paper demonstration (see Fig. 3). In addition, we presented three tag clouds from popular websites (flickr, delicious, last.fm) that all followed a sequential layout with alphabetical sorting. We expected that many participants had already come across this well-known layout – effects caused by a higher familiarity cannot be avoided but somewhat compensated by presenting this type of tag cloud to all participants in advance.

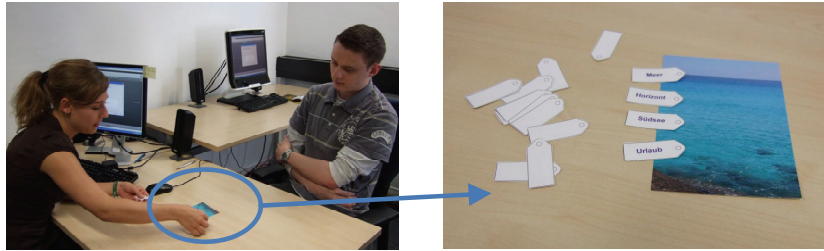


Fig 3. Introduction into tagging and tag clouds by oral explanation and a paper demonstration.

After the introduction, the participants were randomly assigned to one of the three tasks, resulting in 12 participants per task. We applied a Graeco-Latin square design to counterbalance the layout-corporus-combinations in the series of trials for each experimental group. Table 1 shows the resulting four series for task 1. The second and third experimental groups (task 2 and 3) were assigned to corresponding layout-corporus-combinations that also followed a Graeco-Latin square design.

Table 1. Graeco-Latin square design for the first task (group 1): Systematizing tag cloud layout (L) and tag corpus (C) for the four subgroups (SG) and four trials (Tr).

SG	Tr1	Tr2	Tr3	Tr4
1	L1 / C1	L2 / C3	L3 / C4	L4 / C2
2	L2 / C2	L1 / C4	L4 / C3	L3 / C1
3	L3 / C3	L4 / C1	L1 / C2	L2 / C4
4	L4 / C4	L3 / C2	L2 / C1	L1 / C3

The tag clouds were presented on a 17" TFT monitor with a screen resolution of 1280 x 1024 px. They were placed in the middle of a blank screen in an area of 20 x 13.3 cm. Before the presentation of each layout, a short text was displayed explaining the scenario and task. The text of the task was identical within each group; the scenario was adapted to the corresponding tag corpus. The participants should carefully read the task and the scenario before moving on. When one layout was completed by clicking on the tags, the group's task was again displayed along with the next scenario and tag cloud layout. After all four layouts had been presented, the participants filled out a questionnaire in which they chose their preferred layout among others. All four layouts were again shown in the questionnaire to aid recognition. The click times were recorded by the presentation software. Gaze data was captured by an eye tracking system that was embedded in the TFT monitor.

5 Results of Task-Layout-Performance

We measured the task-layout-performance as the time the participants needed to accomplish the tasks. For the first task, we measured the time until the target tag was selected. For the second and third tasks, we calculated the mean value of the time needed for the three selections as indicator for task-layout-performance.

We ran a Kruskal-Wallis test on the data to get an objective ranking on how the layouts perform for each of the three tasks. In addition, we calculated a subjective ranking by counting the study participants' votes for the layouts in the post-test questionnaire. Table 2 shows these values for all three tasks along with the means, medians, minima, and maxima (all in seconds) of the times needed. In the following, we discuss the results for each task in more detail.

Table 2. Performance values of the tag cloud layouts for the 12 participants (N) of each task: Kruskal-Wallis (KW) mean rank, user votes, mean, median, minimum, and maximum (in sec).

Task	Layout	N	KW Mean Rank	User Votes	Mean	Median	Min	Max
1 p=.131 3df $\chi^2=5,6$	sequential	12	22.5	8	8.6	6.5	2.5	18.8
	circular	12	27.8	1	13.6	12.8	3.3	42.3
	clustered	12	30.0	0	14.3	12.2	2.8	33.5
	reference	12	17.8	3	6.9	5.3	2.3	18.7
2 p=.036 3df $\chi^2=8,5$	sequential	12	20.8	1	2.6	1.6	1.0	6.6
	circular	12	18.7	8	2.2	1.8	0.8	4.7
	clustered	12	24.4	3	3.1	2.3	1.0	7.8
	reference	12	34.1	0	4.5	3.4	1.2	9.8
3 p=.239 3df $\chi^2=4,2$	sequential	12	26.5	4	7.3	5.0	1.7	23.0
	circular	12	22.0	3	5.4	4.7	1.9	12.2
	clustered	12	19.4	5	4.9	3.7	1.5	13.2
	reference	12	30.2	0	6.9	6.5	3.0	11.5

5.1 Results for Task 1

For the first task (finding a specific tag), the reference layout (the layout without any weighting of tags) performed best with a value of 17.8 when comparing the mean ranks, followed by the sequential tag cloud (mean rank: 22.5). Since both layouts were alphabetically sorted and hence highly supported scanning for a particular term, we expected these layouts to score well in this task, though the differences were not significant (p=.131). Furthermore, we found some evidence that tags with large font sizes are selected more quickly than small tags in this task. However, we cannot provide reliable values as this effect seemed to be influenced by other visual features, such as a tag's number of characters or its position and neighboring tags in the cloud.

Interestingly, only 3 of this group's 12 participants voted for the reference layout in the questionnaire (two of them have never come in contact with tag clouds before). Although the font sizes did not provide additional value in this task, most participants (8 of 12) preferred the sequential tag cloud layout. It seems that their choice was not only driven by rational factors but largely influenced by aesthetic aspects, since many subjects mentioned the more appealing appearance of the tag clouds compared to the reference layout in the questionnaire.

5.2 Results for Task 2

In the second task (finding the most popular tags), the circular layout showed the significantly best performance results (mean rank of 18.7, $p=.036$). It was also selected as the preferred layout by most participants (8 of 12). Unsurprisingly, the reference layout failed both in click times and user ratings as it does not provide any information about tag popularity.

Interestingly, the two participants that needed the longest time to accomplish the task preferred the clustered layout – it seems as if the circular layout supports the identification of objectively popular tags (i.e., tags with large font sizes) whereas the clustered layout encourages subjective decision making on most important tags (i.e., considering not only a tag's font size but also its position in the clustered layout). However, this interpretation needs further validation.

5.3 Results for Task 3

The results of the third task (finding tags that belong to a given topic) also met our expectations as the clustered tag cloud layout performed best (though not significantly, $p=.239$). However, not all participants seem to recognize its ordering principle; some even selected tags from topic areas that were not asked for in the task. Interestingly, most participants that chose tags in accordance with the task also voted for the clustered tag cloud in the questionnaire.

Again, font size had a strong effect on tag selection, as the selected tags had a mean size of 13.6 pt (median: 25 pt) compared to a mean size of 19.6 pt (median: 18 pt) for all tags, even though the study participants could freely choose any tags that belong to the given topic area, independently of their font size.

6 Analysis of Eye Tracking Data

The analysis of the eye tracking data was performed in an exploratory manner looking for typical patterns in the visual search. We focused on the spatial distribution of fixations over the three tag cloud layouts (the reference layout was not of interest here). All fixations of the subjects' eyes with a minimum duration of 100 ms were considered, using a fixation radius (i.e., the smallest distance that separates fixations) of 50 px. We analyzed the fixation data of the first six seconds of each trial, which represent well the main phase of the search process.

In a first analysis, we divided the whole tag cloud areas into 5 x 5 rectangular, equally sized subareas and counted fixations in these areas. Fig. 4 shows the distribution of fixations over the 25 subareas. In the circular layout, fixations are quite strongly focused on the central part of the tag cloud. They are more dispersed and oriented towards the upper left quadrant in the sequential layout. In the clustered layout, they are more evenly distributed. The lower right attracts less attention in all layouts.

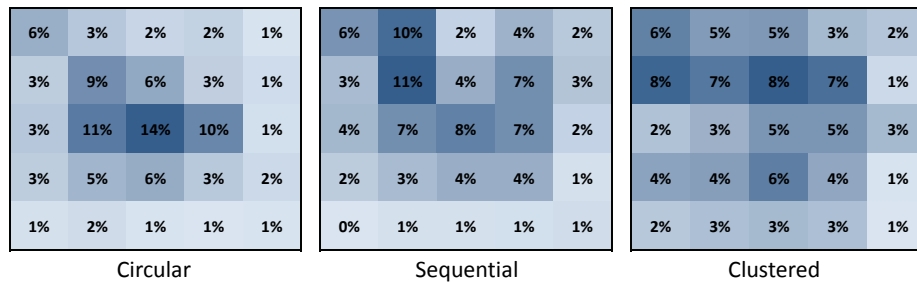


Fig. 4. Distribution of fixations in percent over the 5 x 5 subareas for the three tag cloud layouts. The five-level coloring illustrates the pattern of the distribution.

The uneven distribution of fixations among the quadrants becomes even more apparent when aggregating the 25 subareas into partly overlapping quadrants (see Fig. 5a). Compared to a strict quadrant separation, fixations near the quadrant borders are counted for more than one quadrant in this variant what further strengthens the results in our case. Interestingly, the distribution of fixations over the quadrants varies only marginally among the layouts (see Fig. 5b). In all layouts, the upper left quadrant got the most fixations (32-36%) and the bottom-right the fewest (19-21%). This might be explained by Western reading habits.

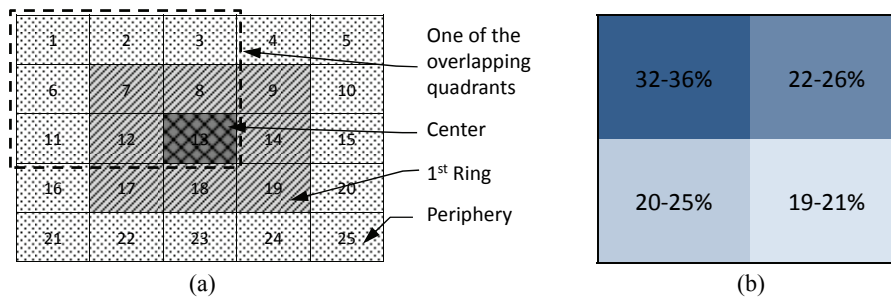


Fig. 5. a) Definition of areas of interest, b) Distribution of fixations over the quadrants

To further analyze the tendency towards the center we aggregated the 25 subareas to three larger, concentric zones as illustrated in Fig. 5a: 1) a center zone consisting of the single rectangle in the middle, 2) a 1st ring of rectangles around the center, and 3) a 2nd ring of rectangles ('periphery'). In this case, we calculated the number of fixations per unit of area (fixation density). The data for the center, 1st ring, and

peripheral zones show that the highest proportion of fixations lies in the central parts of the cloud (center + 1st ring), regardless of the layout (see Fig. 6). For the circular layout, the small center rectangle has even more than double the fixation density than the surrounding 1st ring, indicating a very strong central focus.

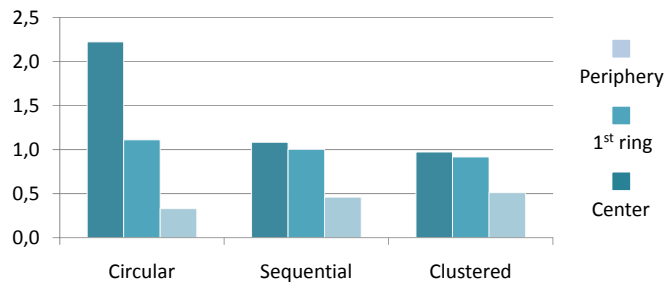


Fig. 6. Fixation density for the central-to-peripheral zones

In order to gain insight into the temporal variation of the gaze focus, we divided the overall time interval analyzed (sec 0-6) into three subintervals: sec 0-1 (period 1), sec 1-3 (period 2), and sec 3-6 (period 3). We chose three periods of increasing length since we were interested in contrasting a short initial phase indicative of a subject's orientation at the beginning of the trial with the gaze focus in later intervals where the visual search becomes more expansive and covers larger areas. Fig. 7 shows the gaze distribution over the center-to-periphery zones for the three time periods. As already observed in the overall gaze distribution, the circular layout generates the highest central orientation, which is slightly decreasing over time for the benefit of the 1st ring and peripheral zones. The change in distribution is less pronounced for the sequential layout. In the clustered layout, the decrease in central orientation is strongest. A possible explanation is that the gaze tends to stay in the target cluster once the thematic relation with it has been recognized.

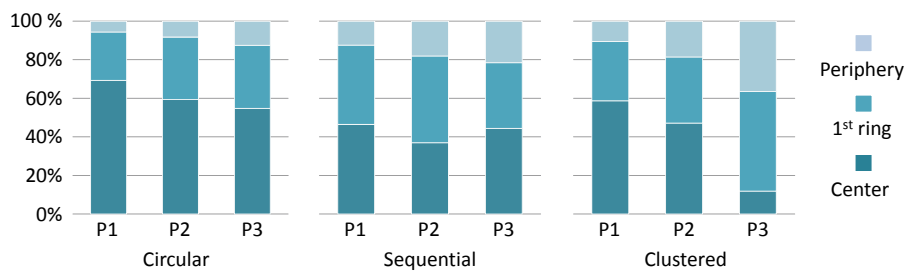


Fig. 7: Distribution of fixation density in the center, 1st ring and periphery zones for the three time periods.

7 Conclusions

The results of our comparative study on tag cloud layouts show clearly that there exists no single best way to arrange weighted terms in a cloud – as common in interaction design, the optimal solution depends strongly on the specific user goals and intentions of the designer. However, regarding the layout classes and user goals investigated in our experiment, the following task-layout-matchings are supported by our results:

1. Finding a specific tag: Sequential layout with alphabetical sorting.
2. Finding the most popular tags: Circular layout with decreasing popularity.
3. Finding tags that belong to a certain topic: Thematically clustered layout.

Generalizing from the layouts and tasks, we got some results that are largely in line with the findings and assumptions made in related work. Table 3 summarizes the findings by contrasting the general results of our study with the results of related work (cp. Sec. 2): As already indicated in previous studies [2, 5, 9], we also found that tags with large font sizes seem to ‘pop out’ to the viewers – they were on average identified more quickly in the first task and more often selected in the third. Although tag clouds perform worse than unweighted lists on average [5, 8] they seem to accelerate the identification of very popular tags.

Table 3. Summarization of the research findings on tag cloud perception and performance contrasting related work and our own study (encompassing both quantitative and qualitative observations)

Abbreviation	Description & literature source	Evaluation based on our results
Tag size	Large tags attract more user attention than small tags [2, 9] and are found more quickly [5].	Basically, our results support these findings, though they also indicate that further properties, such as the number of characters, the position of a tag in the cloud, or neighboring tags, influence this effect (cp. [2]).
Scanning	Users scan rather than read tag clouds [2, 5].	The eye tracking data support this finding. However, in contrast to the assumption of [5], no general scanning direction could be identified.
Centering	Tags in the middle of the cloud attract more user attention than tags near the borders [2].	This seems to be true for most kinds of tag clouds though the layout can increase (e.g., circular) or decrease (e.g., clustered) this effect.
Position	Tags in the upper left quadrant are better recalled [9] and are found more quickly [2].	The eye tracking data show that the upper left quadrant receives the most fixations in all layouts.
Layout	The layout of a tag cloud influences its perception [9].	This finding is strongly supported by our eye tracking data.
Exploration	Tag clouds provide suboptimal support when searching for specific tags [5, 8, 10].	In general, we support this finding. However, very popular tags can be found comparatively quickly.

When aggregating the eye tracking data of all layouts and tasks, we observed a tendency towards the center: Tags in the middle of the cloud attracted more user attention than tags near the borders. Also, an effect of position could be observed in the eye tracking data: The upper-left quadrant received more fixations than the others. Thus, a tag's general visibility is largely influenced by its position in the tag cloud, which might be used by designers to direct users' attention to certain tags.

The feedback and ratings of the study participants suggest that fun and aesthetic aspects largely affect the user's interaction with tag clouds. Participants partly preferred layouts that did not yield the best performance. Thus, such aspects should be included in any design decision and usability evaluation of tag clouds. A holistic approach must furthermore consider the general nature of tags and the characteristics of tagging systems [4] that were not within the focus of this paper. As shown by our findings, interface designers must also and foremost consider the possible user tasks before deciding for a certain type of tag cloud or using tag clouds at all.

References

- [1] Ambrose, G., Harris, P.: Grids. Ava Acedemia, Lausanne (2008)
- [2] Bateman, S., Gutwin, C., Nacenta, M.: Seeing Things in the Clouds: The Effect of Visual Features on Tag Cloud Selections. In: Proc. of the 19th ACM conference on Hypertext and Hypermedia, pp. 193–202. ACM Press, New York (2008)
- [3] Clark, J.: Clustered Word Clouds (2008), <http://www.neoformix.com/2008/ClusteredWordClouds.html> (retrieved on 2009-01-30)
- [4] Golder, S., Huberman, B.A.: Usage Patterns of Collaborative Tagging Systems. *Journal of Information Science* 32, 2, 198–208 (2006)
- [5] Halvey, M., Keane, M.T.: An Assessment of Tag Presentation Techniques. In: Proc. of the 16th International Conference on World Wide Web, pp. 1313–1314. ACM Press, New York (2007)
- [6] Hassan-Montero, Y., Herrero-Solana, V.: Improving Tag-Clouds as Visual Information Retrieval Interfaces. In: Proc. of International Conference on Multidisciplinary Information Sciences and Technologies. INSTAC, Mérida (2006)
- [7] Hearst, M.A., Rosner, D.: Tag Clouds: Data Analysis Tool or Social Signaller? In: Proc. of the 41st Annual Hawaii International Conference on System Sciences. IEEE Press, New York (2008)
- [8] Kuo, B.Y.-L., Hentrich, T., Good, B.M., Wilkinson, M.D.: Tag Clouds for Summarizing Web Search Results. In: Proc. of the 16th International Conference on World Wide Web, pp. 1203–1204. ACM Press, New York (2007)
- [9] Rivadeneira, A.W., Gruen, D.M., Muller, M.J., Millen, D.R.: Getting our Head in the Clouds: Toward Evaluation Studies of Tagclouds. In: Proc. of the SIGCHI Conference on Human Factors in Computing Systems, pp. 995–998. ACM Press, New York (2007)
- [10] Sinclair, J., Cardew-Hall, M.: The Folksonomy Tag Cloud: When is it Useful? *Journal of Information Science* 34, 1, 15–29 (2008)
- [11] Viégas, F.B., Wattenberg, M.: Tag Clouds and the Case for Vernacular Visualization. *Interactions* 15, 1, 49–52 (2008)
- [12] Popular Tags on Flickr, <http://www.flickr.com/photos/tags/> (retrieved on 2009-01-30)