

Visualizing Tags with Spatiotemporal References

Dinh Quyen Nguyen, Christian Tominski, Heidrun Schumann
Institute for Computer Science
University of Rostock
Rostock, Germany
{nguyen|ct|schumann}@informatik.uni-rostock.de

Tuan Anh Ta
John von Neumann Institute
Vietnam National University - HCMC
Ho Chi Minh City, Vietnam
tuan.ta@jvn.edu.vn

Abstract—Nowadays, a great amount of data is created and distributed on the Internet. Tagging has become common practice to structure these data for easy access. Often the data and the associated tags contain spatial and temporal information.

In this paper, we develop general design strategies for visualizing spatially and temporally referenced tags similar to tag clouds on maps. Temporal information of tags is encoded through the visual appearance of text or through additional visual artifacts associated with the tags, whereas the location of tags on a map illustrates the spatial references. We demonstrate our solution based on an interactive visualization prototype for the exploration of both spatial and temporal references of Flickr tags.

Index Terms—Visualization, Tagging, Spatiotemporal Data, Flickr, Tag Clouds

I. INTRODUCTION

Keywords from an article or a book offer a special kind of linguistic summary for text. Beyond that, there is the concept of tagging, where any kind of media (e.g., music, movies, images) can be associated with textual tags in order to create such linguistic summaries.

Keywords or tags can be represented visually as tag clouds. In a tag cloud, not only the words are meaningful, but the appearance of the words (e.g., size, color, etc.) can be used to convey information, for example the importance or frequency of a word. Nowadays, tag clouds are commonplace on many websites because they efficiently communicate an overview of the site’s content. In this work, we are interested in this topic of text visualization, but we consider two additional aspects: our tags are associated with geo-spatial and temporal information.

A photo hosting website such as Flickr is a prominent application scenario for spatiotemporal tags. When a photo is taken or uploaded it is automatically equipped with a time stamp. Additionally, many photos are also outfitted with spatial information such as precise geo-positions (i.e., latitude and longitude) or coarser areal information (e.g., west coast or city of Munich). Last but not least, the photographer or the user community may tag a photo with keywords (simply called tags) that describe the photo’s content.

In turn, this means that tags are no longer just abstract linguistic objects. Instead, a tag carries meaning and embeds this meaning into a spatiotemporal frame of reference. Given such spatiotemporally referenced tags, a user might want to find out where certain tags occur, or how the tags are linked across different areas. He or she might also be interested in exploring the tags with regard to a specific time point or time

interval of interest, or with regard to a specific temporal pattern (e.g., consecutive Sundays of a month). Examining how tags have evolved over time is also an interesting task.

While common tag clouds are useful for visually communicating tags, they cannot convey both spatial and temporal aspects. Yet, there are tag clouds that are capable of showing either the spatial or the temporal context of tags. Yahoo! Tag Maps [1] and Taggram [2] are examples of existing techniques for the visualization of tags on maps. On the other hand, Yahoo! Taglines [3] or Twitscoop [4], amongst some others, provide ways for the visual representation of tags over time.

We develop a first solution for the visualization of both spatial and temporal dependencies of tags. For the spatial component of the data, we utilize the concept of Taggram. In order to allow users to recognize temporal information, we enhance the visual representation of tags in Taggram. In accordance to McEachren’s list of aspects of time [5], we aim to support users in answering the following set of questions:

- Q1: Does a specific tag exist at a particular time?
- Q2: When in time does a tag occur?
- Q3: How old is a tag at a specific time point (if existing)?
- Q4: For how long does a tag exist?
- Q5: How often is a tag used?
- Q6: How frequent or strong does a tag change over time?
- Q7: Is there any special temporal relationship (e.g., co-occurrence) between some tags in the cloud?

Our goal is to develop general designs for representing the different temporal aspects of tags. The suggested solution is two-fold: (i) time is encoded through the visual appearance of text (e.g., font size, text color, orientation, transparency, shape), and (ii) time is represented by additional visual artifacts associated with the tags.

The paper is organized as follows. In Section II, we take a brief look at related work and the fundamentals of Taggram. Section III introduces our general design strategies for the visualization of tags associated with temporal references. In Section IV, we present concrete design examples for the visualization of spatiotemporal Flickr tags developed based on our novel strategies and Taggram. Section V describes the interaction facilities provided to support the visual exploration. We briefly describe implementation aspects and report on preliminary user feedbacks in Section VI. Finally, Section VII concludes the paper.

II. RELATED WORK & FUNDAMENTALS

Because there is hardly any visualization that displays spatial AND temporal references of tags, we summarize previous work related to either space OR time. We also briefly explain Taggram, which is the basis for the spatial mapping.

A. Related Work

Tag clouds have become very popular on the Internet over the past few years. The aim of this kind of visual representation is to provide users with a simple and expressive overview of a larger body of content (e.g., a set of online news articles). Commonly, the basic technique is straight forward: display a set of words on the screen, where font size, color, and orientation visually encode some related information (e.g., frequency of a word). However, more information can be visually encoded with tag clouds:

1) *Tag Clouds and Space*: In the theory of perception, shape is a prominent visual variable for information acquisition [6]. Therefore, the shape of an area linking with a tag cloud can be used to represent contextual information of the tags (without losing the validity of those tags). For example, in DocuBurst [7], the hierarchical structures of the interrelated words are visualized as a polar treemap that can be explored interactively.

Many tag cloud visualization approaches focus on the question how to spatially layout (aka spatialize) words on the display? It is also the basic idea behind text typography, a special kind of information graphics [8]. Interestingly, it is also relevant to the problem of using tag clouds in geospatial data visualization. Taggram [2] is a recently developed technique that addresses this problem. However, Taggram only concerns the matter of spatial references of tags; temporal features have not been taken into account.

2) *Tag Clouds and Time*: Because online communities are usually dynamic, there are a number of implementations that integrate the one or the other aspect of time into tag clouds.

The tool Twitscoop [4] uses the slideshow approach. It updates the visualization every minute to reflect the currently most important keywords in Twitter. Yahoo! Taglines, for instance, uses animation to visualize changes of topics over time [3]. The user can choose from two animation schemes (river or waterfall), where time goes from left to right and top to bottom, respectively. If a tag cloud spans only a small number of time steps or if just the latest snapshot of the tag cloud is relevant, a dynamic visual representation (slide show or animation) is suitable, because the temporal changes are usually easily comprehensible.

In order to facilitate more detailed visual analysis, additional visual cues can be used. Stefaner uses animation as well, but additionally encodes information about the "age" of tags [9]. Cui et al. [10] use color to differentiate newer and older tags in the cloud. But comparisons of different time steps or different tags are still difficult to conduct. In this case, combining tag clouds with additional views (e.g., a temporal bar chart as in [10]) is one option to solve this problem. A similar implementation is offered by Cloudalicious [11].

While most approaches rely on basic techniques and show only one snapshot of a tag cloud at a time, only few attempts have been made to visualize data of multiple time steps concurrently or to support the analysis of temporal aspects such as frequency or rate of change. One such approach is SparkClouds [12], which combine spark lines with tags. But still, visualizing time-varying tag clouds remains an interesting research topic.

B. Taggram

We develop our solutions for the visual representation of spatiotemporally referenced tags based on Taggram [2]. Taggram is a flexible technique that deals with the spatialization of geo-referenced tags inside geographic regions. The key characteristics of this approach are:

- Tags are placed on a map according to the tags' spatial references. Regions of arbitrary shape are used as visual containers for the spatialization of tags. Tags that belong to a geographical region are placed inside that region in alphabetical order, following the main vertical axis of the region.
- Information related to tags can be encoded through text attributes (e.g., color or transparency) and also through additional visual artifacts that can be superimposed on the map. Note that size – conventionally a basic visual attribute of tag clouds – is used differently in Taggram: Tags that are close to the center of the region are enlarged, and tags with increasing distance to the center become smaller. The rationale behind this design is that normally the user prefers to explore tags at the center of a tag cloud [13].
- Because the number of tags in the data is usually larger than the number of tags that fit the available display space, tags can be dynamically scrolled inside the geographic region on demand. This way, users can stay focused on the region's center while exploring the alphabetically ordered tags.

With these characteristics, Taggram represents an excellent basis for our development. We can conveniently add more visual artifacts to the tags to represent their temporal aspects. In the next section, we will present general design strategies for the visualization of time-oriented tags.

III. GENERAL DESIGNS FOR THE VISUALIZATION OF TAGS WITH TEMPORAL REFERENCES

According to Aigner et al. [14], time-oriented data can exhibit a number of different characteristics, which leads to a variety of potential visualization designs. So, there are various ways to visualize temporally referenced tags. In our work, we focus on adapting the visual representations of the tags themselves, in order to arrive at designs that are suitable to answer the questions Q1 to Q7 listed in Section I. Note that the fact alone that a tag appears in a tag cloud enables users to recognize existence of the tag (Q1).



Fig. 1. Encoding temporal aspects with text appearance.

Next, we discuss two general design strategies for visualizing the temporal aspects: (i) by the visual appearance of the tags, and (ii) by additional visual artifacts associated with tags.

A. Temporal Aspects Mapped to Tag Appearance

Basic visual attributes such as color, size, or orientation are normally employed to visually encode information associated with tags. Visual attributes are also applicable to encode temporal aspects of tags. For example, one can use hue, saturation, or brightness to differentiate tags that have appeared at various points in time. Fig. 1(a) illustrates the encoding with brightness: more recent tags are darker and older tags are brighter. Similarly, text size can indicate how recent a word is (see Fig. 1(b)). For both encodings, the assumption is that recent tags are more important.

Provided that appropriate legends are displayed, such visual encodings allow users to discern a tag’s age (Q3), and thus to estimate a tag’s location in time (Q2). Conceptually, we are now able to visualize tags and a single data value per tag (e.g., age). But this is not enough for explicitly visualizing multiple time steps concurrently or for supporting the analysis of temporal aspects such as the frequency or rate of change. In order to arrive at conclusions for the questions Q4 to Q7, we need to consider other visual encodings.

One option is to modify the text rendering along the tag. The basic idea is to map individual time axes along the tags’ horizontal orientation. That is, the beginning of a tag corresponds to the first time step and the end of a tag corresponds to the last time step. Given this mapping, we can vary the rendering along the tag. In Fig. 1(c), we show tags that have been distorted to visualize how frequent the tags appear over time. For example, the letters of “always” are larger throughout, indicating that the tag has been important at all points in time. On the other hand, the letters of “recently” increase in size towards the end of the tag, indicating a tag of growing importance. Similarly, one can vary other rendering parameters such as transparency as shown in Fig. 1(d).

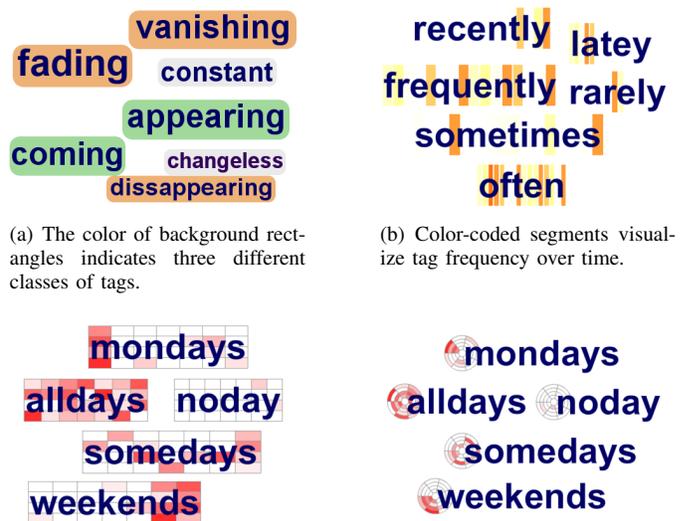


Fig. 2. Encoding temporal aspects with additional visual artifacts.

Provided that there are no uncommonly short or long tags, these encodings along the tags’ horizontal axes are suitable for indicating how tags have developed over time. For a more detailed exploration, however, varying basic visual attributes alone is not sufficient. Therefore, we suggest using additional visual artifacts as described next.

B. Temporal Aspects Mapped to Additional Visual Artifacts

By additional visual artifacts we mean enriching tags with additional graphical primitives. The main advantage of this solution, in comparison to the designs of the previous paragraphs, is its flexibility: By encoding temporal information with the visual attributes of the additional graphical primitives, we can visualize more complex temporal information. A positive side effect is that basic visual features of the tags (i.e., color, size, etc.) are kept for the traditional encoding of tag importance.

Fig. 2(a) shows a first simple example where a background rectangle has been added to each tag. The rectangle’s color is used to indicate three groups of tags with different temporal behavior: appearing tags are green, disappearing tags are orange, and constant tags are gray.

More complex designs can communicate further information. Again, the basis is to construct an appropriate mapping of the time axis. For example, one can use color-coded segments along the horizontal extent of a tag to show the frequency of the tag over time (see Fig. 2(b)). This design is suited to visualize time in a linear fashion. Cyclic recurrence of tags can be made comprehensible with the help of table-based calendars (see Fig. 2(c)) or glyph-based cyclic artifacts (see Fig. 2(d)). These designs can also help in identifying any regular temporal patterns (e.g., tags occur only at particular weekdays).

With the aforementioned designs, the remaining questions Q4 to Q7 can be answered: The duration of tag existence, the frequency of tags, as well as the rate of change appear vividly in Fig. 2(b)-2(c), and temporal patterns are recognizable in Fig. 2(c)-2(d). Moreover, we can visualize time points together with time intervals, show linear time together with cyclic patterns, and switch to alternative perspectives, in order to explore the various temporal aspects of tags.

Further options are to compose different designs to show different temporal aspects as hybrid visual representations. However, care has to be taken to avoid cluttering the tags and to maintain the tag cloud's legibility. In this regard, an important aspect is to support user interactions such as zooming into details, highlighting tags or time points of interest, or interactive adjustment of the visual encoding.

The next consequent step is to extend these generic designs to address a more concrete visualization scenario.

IV. VISUALIZATION OF SPATIOTEMPORAL FLICKR TAGS

In contrast to the generic designs presented before, we now address the visualization of spatiotemporal tags as provided by Flickr. Furthermore, we take into account the fact that a user is usually interested in a specific time point, the so-called time of interest (TOI). In Fig. 3, for example, Fri, 8/20/2010 has been selected as TOI. The TOI divides the time domain into three parts: the time before the TOI, the TOI itself, and the time after the TOI. Note that the before and after parts might be empty, if the TOI is the beginning or the end of the time domain.

Given this additional TOI concept, we have to develop dedicated visual encodings, for which we again utilize (i) the visual appearance of tags and (ii) additional visual artifacts. Again, it is our goal to have at least a set of techniques that covers the questions Q1 to Q7.

A. Adjusting the Visual Appearance of Tags

The two visual designs that we present next are based on masking the tags. The designs operate on two different semantic levels. For the first design, we focus on qualitative statements regarding the existence of tags in the time domain. In the second design, we extend to quantitative assessments regarding the frequency of tags over time.

1) *Encoding for Qualitative Statements:* With regard to the TOI concept and the existence of tags, our design must be capable of communicating seven cases. The seven cases can be represented visually by color gradients, where red stands for existence, and gray indicates non-existence of a tag. The following list summarizes the seven cases of tag existence and the corresponding visual encoding: A tag exists

- | | |
|----------------------------------|--|
| 1) only at the TOI: | |
| 2) only before the TOI: | |
| 3) only after the TOI: | |
| 4) at the TOI and before: | |
| 5) at the TOI and after: | |
| 6) before and after the TOI: | |
| 7) before, after and at the TOI: | |



(a) Tags masked with gradient fills visualize tag existence.



(b) Tags masked with color-coded shapes visualize tag frequency over time.

Fig. 3. Visualization of Flickr tags on maps with temporal aspects represented by text appearance. (Recall that Taggram shows tags in alphabetical order and varies text size to attract the viewer's attention to the center of a region.)

To establish a visual association of a tag with one of the seven cases, the tag is masked with the gradient fill. This way, each tag is rendered according to the case it belongs to. Fig. 3(a) illustrates this with a sample data set extracted from Flickr: (1) tag “magical”, (2) tag “messe”, (3) tag “launsbach”, (4) tag “licht”, (5) tag “lowersaxony”, (6) tag “nikon”, and (7) tag “museum”. Note that the combinatorial eighth case (i.e., the tag does not exist at all) must not be handled, because in that case the tag does not appear at all.

From the encoding presented here, the user can derive only qualitative information about tag existence (Q1 and Q2). How quantitative values such as tag frequency can be encoded will be shown next.

2) *Encoding for Quantitative Assessment:* Now our goal is to visualize data values associated with tags such as tag frequency (i.e., the number of occurrences of a tag over time per geographical region). What we need is a visual mapping of the time domain that is capable of showing quantitative values per time point.

To this end, we create series of visual segments (one for each time point) along the horizontal orientation of tags. We continue using the idea of the TOI and use different shapes for the segments. A diamond shape \diamond is used to represent the TOI, whereas time points before and after the TOI are

represented as arrow shapes \Leftarrow and \Rightarrow , respectively. We use colors from the ColorBrewer [15] to color-code each shape according to the tag’s frequency at the corresponding time point. As tag frequencies might be distributed unevenly over time, additional color mapping concepts from [16] can be applied to redistribute colors for more effective visualization. For those time points where the frequency is zero, we use a hueless color to clearly indicate the non-existence of that tag at that time point.

These mapping steps result in visual representations like  for each tag. Again, we mask the tags to integrate the visualization of the time-dependent frequency values with the visual appearance of tags. Fig. 3(b) illustrates such masked tags for a part of Germany for a time range of 8 days in August 2010. This design is suited to find answers to the questions Q3 to Q6. However, if the time domain is larger, differentiating the individual time points may become difficult. In such cases, one could exploit the hierarchical structure of time and represent aggregated frequencies of weeks or months, rather than those of individual days. Another option is to associate additional visual artifacts with tags.

B. Additional Visual Artifacts Associated with Tags

As indicated earlier, using the visual appearance of tags alone might not be sufficient for more complex questions or larger time domains. Therefore, we now consider adding visual artifacts to the tags on the map. Each artifact can then be used to visualize the temporal aspects in more detail.

1) *Bar Charts Representing Temporal Developments:* In this first design, we use bar charts as additional visual artifacts to visualize the existence and frequency of tags over a period of time (Q1 - Q6). We chose bar charts because they are an accepted method for visualizing time-dependent data and because they are easy to interpret.

The design of a bar chart per tag is straightforward: For each time point, we create a rectangular bar to show if the tag existed and how frequent it was. We add a small red dot to indicate the TOI. Color and height of bars are used to show tag frequency at individual points in time. The color-coding is the same as described in the previous section. Using a dual encoding with color and height has the positive side effect that we do not have to deal with the special case of non-existence, because in that case the bar has zero height and is invisible anyway.

The bar chart design is illustrated in Fig. 4(a) with a list of Flickr tags for the period of 8/9 to 8/29 in Germany, where Sat, 8/21/2010 has been selected as the TOI. From the bar charts in this visual representation, the user can easily see how tags appeared and if there are any linear temporal patterns. However, bar charts are not suited to find more complex and possibly reoccurring patterns.

2) *Calendar Tables Representing Temporal Patterns:* Color-coded calendar tables, whose horizontal and vertical axes represent independent levels of time, are promising alternatives. For example, a common design is to show weeks as rows of a table, which implies that columns represent



(a) Color and height of bars visualize tag existence and tag frequency.



(b) Color of the cells of calendar tables visualize tag frequency.

Fig. 4. Temporal aspects of Flickr tags visualized by additional visual artifacts. (Recall that Taggram shows tags in alphabetical order and varies text size to attract the viewer’s attention to the center of a region.)

individual weekdays. Each table cell’s color indicates tag frequency (or any other attribute users might be interested in).

The user can now spot more complex patterns (Q7), for example, tags that frequently reoccur only on particular week days. At the same time linear trends are discernable, for example, if the frequency increases from one week to the next. Several interesting patterns can be seen in the sample Flickr data visualized in Fig. 4(b) The tag “2010” is significant throughout, but the highest frequencies are on weekends (last two columns of the table). Quite nicely one can see that the event “wiesenhöhne” appears in the middle of the first week (top row of the table) and disappears toward the weekend.

C. Discussion

The two basic design strategies (i.e., visual appearance of tags vs. additional visual artifacts) that we suggested have advantages and disadvantages alike. For both strategies the aim is to visualize additional temporal aspects while maintaining tag legibility. By changing only the appearance of tags, our options for encoding temporal information are limited, and thus only simple questions are supported. On the other hand, tag legibility is only marginally affected (provided that appropriate colors are used). Moreover, the encoding of a tag and its associated temporal aspects is overlap-free.

This is different when using additional visual artifacts. In this case, there can be significant overlap which could occlude possibly important temporal information, for example, when a letter occludes the cell with the highest frequency in a calendar table. This is the price to pay for the additional options that visual artifacts offer for the visualization.

We cannot tell which of the two strategies is best. But most likely there is no definite answer to this question at all. Rather a suitable solution must be selected depending on the data, the users' tasks, and their preferences. Therefore, interactive selection of encodings and their adjustment are important.

V. USER INTERACTION

The visual exploration of spatially and temporally referenced tags involves various aspects, and therefore requires sufficient means for user interaction. On the one hand, the user must be enabled to visit different places in space and to select different time ranges from the data. On the other hand, the visual encoding must be interactively adjustable to the users' needs and preferences.

We support these tasks with a number of interaction techniques. The users can zoom and pan to any place on the map, they can switch to their favorite visualization mode, and they can focus on tags with relevant temporal developments.

To make the different options for interaction easily accessible, we provide a custom-made user interface component, which borrows from the idea of floating menus. Fig. 5(a) shows the main menu with options, including home, settings, color adjustment, map display, photos, and tag selection. When the user clicks the tag selection, the component switches its interface accordingly. Fig. 5(b) shows this interface, which allows users to select the time range for which tags are to be displayed, the time point of interest, as well as the visualization mode (i.e., the different encodings introduced earlier).

Besides presenting tags with temporal dependencies on maps, we also support the users in interacting with the tags for further information exploration as described next.

One important task is to highlight the concrete locations in space where tags originate from. This is necessary because the Flickr database aggregates the spatial information of tags to larger regions. Our goal is to allow users to reestablish the connection of tags to specific locations.



(a) Main menu. (b) Time and encoding selection.

Fig. 5. Controller interface.

To this end, we mark the concrete positions of one selected tag on the map as shown in Fig. 6(a). However, indiscriminately showing all locations could clutter the map with too many markers. Therefore, we restrict the highlighting to one selected time point, usually the TOI. If the visual encoding allows the identification of individual time points (e.g., segment-based appearance of tags, bar chart, or calendar table), the user can choose to select alternative time points for highlighting of locations simply by hovering the time points. In that case, the selected time point is additionally labeled with a tooltip containing date information (see Fig. 6(a)).

Similarly to highlighting spatial aspects, it makes sense to support interactive highlighting of temporal aspects. Identifying tags that co-occur is an interesting task (Q7) because it allows users to derive higher-level information from spatiotemporal tags.

Therefore, we allow users to select further time points in addition to the TOI. The visualization is then adapted so as to highlight those tags that co-occur on all selected time points (which is actually done by dimming the tags that do not co-occur). Fig. 6(b) demonstrates that this interactive highlighting can reveal facts in the data: It appears that there was a "lasershow" in "lowersaxony" in August 2010.



(a) Markers highlight concrete geo-positions for a selected tag.



(b) Selecting multiple time points reveals tag co-occurrence.

Fig. 6. Options for interacting with time-dependent tags on maps.

VI. IMPLEMENTATION AND PRELIMINARY USER FEEDBACKS

We implemented a web-based visualization application to illustrate the concepts introduced in the previous sections based on the available Flickr tags dataset. While the visualization part in our implementation exploits the graphics and interaction capabilities of Flash and ActionScript, the data part is based on MySQL and the Flickr web services APIs.

The Flickr APIs support queries for the relations of tags and geographical regions of the map using the services `flickr.places.tagsForPlace` and `flickr.places.placesForTags`. Both calls expect a range of dates as input and return lists of tags or places as output, respectively.

The corresponding data is as follows. At a particular date, a tag can be associated with various geo-coordinates inside a geographical region (normally concerning many photos). Because in Flickr a geographical region can be indicated through various levels of geographical abstraction (i.e., place types: locality, city, state, etc.), a tag listed for one place can be listed for other places at other administrative levels as well. For example, a tag listed for Berlin (level: city) may also be listed for Germany (level: country). We utilize this for semantic zooming as users zoom into particular regions of the map.

For a place (e.g. a city) and a given date, there exists a list of tags. Consequently, for multiple dates, there are multiple lists of tags, one for each date. Because our visualization client requires efficient access and iteration over multiple dates and geographical regions, we reorganize the retrieved query results in our own database, which consists of multiple lists of tags indexed by date and geographical region.

This data backend drives the efficient visualization of Flickr tags with spatial and temporal references in any Flash-enabled Web browser. Users can access the data for different places by navigating the Taggram-based map and different time ranges can be visited by using the query interface of the controller menu.

Using our implementation, we conducted a small informal test to collect first feedbacks about our visualization designs. Seven students (aged 22-26, 2 females, familiar with Tag Clouds and Yahoo! Flickr) have been asked to use our tool. As this was not a formal evaluation, we did not focus on any particular visualization task or question, but rather we were interested in the general acceptance of our designs for the visual exploration of Flickr tags.

In summary, six out of seven participants liked the fact that they can explore tag clouds concerning both space and time. They were strongly interested in the interactive visualization that highlights the co-occurrence of Flickr tags on the map for selected dates. Additionally, the exploration of temporal patterns (e.g., weekdays in a calendar-based artifact) got special attention.

On the other hand, there were also negative feedbacks. Firstly, adding visual encodings to tags was experienced as

interfering with the legibility of the tags, which we expected to a certain degree. Secondly, some users asked for additional querying mechanisms, in particular a textual search for specific tags of interest.

VII. SUMMARY AND FUTURE WORK

In this paper, we have addressed the visualization of tags with both spatial and temporal references. We apply Taggram to visualize spatial dependencies of tags on a geographical map. Our design strategy for visualizing temporal aspects of tags has been twofold: employ (i) the visual appearance of text or (ii) additional visual artifacts associated with the tags. For both strategies, we described generic designs and developed visualization examples that support the users in interactively exploring spatiotemporal Flickr tags of user-selected time ranges on navigable maps (see Fig. 7). Using our application, the user can accomplish a number of visualization tasks, including identifying dates for which tags exist, finding local trends in the development of tags, and understanding tag co-occurrence over time.

Through this development, we have illustrated that incorporating additional visual cues into the representation of the raw data (i.e., the tags) is a sensible solution to account for additional spatial and temporal references of the data.

Still, many tasks are open for future work. In the context of visualizing temporal dependencies of tags, a more general formal study has to be carried out in order to identify the most suitable combination of visual variables (shape, size, value, grain, hue, orientation, and position). This also includes a comparison of our approaches with existing spatially and/or temporally referenced text visualization techniques. Another issue is that our development mainly focuses on the simple model of linearly ordered time. However, as temporal data can be multifaceted [14], future work may address other innovations for more complex and multi-perspective temporal data. Similarly, one can think of ways to extend our univariate visualizations to cope with multivariate data associated with tags.

ACKNOWLEDGMENT

This work is partly funded by the KAAD. Some of the colors that appear in this work are taken from ColorBrewer. Some of the icons that appear in this work are by Yusuke Kamiyamane.

REFERENCES

- [1] S. Ahern, M. Naaman, R. Nair, and J. H.-I. Yang, "World Explorer: Visualizing Aggregate Data from Unstructured Text in Geo-Referenced Collections," in *Proc. of the ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL)*. New York, NY, USA: ACM, 2007, pp. 1–10. [Online]. Available: <http://doi.acm.org/10.1145/1255175.1255177>
- [2] D. Q. Nguyen and H. Schumann, "Taggram: Exploring Geo-data on Maps through a Tag Cloud-Based Visualization," in *Proc. of the International Conference on Information Visualisation (IV)*. IEEE Computer Society, 2010, pp. 322–328.
- [3] M. Dubinko, R. Kumar, J. Magnani, J. Novak, P. Raghavan, and A. Tomkins, "Visualizing Tags Over Time," in *Proc. of the International Conference on World Wide Web (WWW)*. New York, NY, USA: ACM, 2006, pp. 193–202. [Online]. Available: <http://doi.acm.org/10.1145/1135777.1135810>



Fig. 7. Visualization of spatiotemporal tags on maps using different encodings.

- [4] "Twitscoop," retrieved Mar, 2011. [Online]. Available: <http://www.twitscoop.com/>
- [5] A. M. McEachren, *How Maps Work: Representation, Visualization, and Design*. New York, NY, USA: Guilford Press, 1995.
- [6] J. Bertin, *Semiology of Graphics: Diagrams, Networks, Maps*. Redlands, CA, USA: Esri Press, 2011.
- [7] C. Collins, S. Carpendale, and G. Penn, "DocuBurst: Visualizing Document Content using Language Structure," *Computer Graphics Forum*, vol. 28, no. 3, pp. 1039–1046, 2009.
- [8] J. Clark, "Twitter List Profile Clouds," retrieved Mar, 2011. [Online]. Available: <http://neoformix.com/2009/TwitterListClouds.html>
- [9] M. Stefaner, "Visual Tools for the Socio-semantic Web," Master's thesis, University of Applied Sciences Potsdam, 2007.
- [10] W. Cui, Y. Wu, S. Liu, F. Wei, M. X. Zhou, and H. Qu, "Context Preserving Dynamic Word Cloud Visualization," in *Proc. of the IEEE Pacific Visualization Symposium (PacificVis)*. IEEE Computer Society, 2010, pp. 121–128.
- [11] T. Russell, "Cloudalicious: Folksonomy Over Time," in *Proc. of the ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. ACM, 2006, p. 364.
- [12] B. Lee, N. H. Riche, A. K. Karlson, and S. Carpendale, "SparkClouds: Visualizing Trends in Tag Clouds," *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, no. 6, pp. 1182–1189, 2010.
- [13] S. Lohmann, J. Ziegler, and L. Tetzlaff, "Comparison of Tag Cloud Layouts: Task-Related Performance and Visual Exploration," in *Proc. of the IFIP TC13 Conference Human-Computer Interaction (INTERACT)*. Springer, 2009, pp. 392–404.
- [14] W. Aigner, S. Miksch, W. Müller, H. Schumann, and C. Tominski, "Visual Methods for Analyzing Time-Oriented Data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 1, pp. 47–60, 2008.
- [15] "ColorBrewer," retrieved Mar, 2011. [Online]. Available: <http://www.colorbrewer.org/>
- [16] C. Tominski, G. Fuchs, and H. Schumann, "Task-Driven Color Coding," in *Proc. of the International Conference on Information Visualisation (IV)*. IEEE Computer Society, 2008, pp. 373–380.