Representing and Visualizing Folksonomies as Graphs – A Reference Model

Steffen Lohmann
University of Stuttgart
Universitätsstraße 38
Stuttgart, Germany
steffen.lohmann@vis.uni-stuttgart.de

Paloma Díaz
Universidad Carlos III de Madrid
Avenida de la Universidad 30
Leganés, Spain
pdp@inf.uc3m.es

ABSTRACT

We present a reference model for the representation and visualization of folksonomies as graphs. We discuss the formal representation of folksonomies and derive a hypergraph structure that is core to any visualization approach. We split this hypergraph into subgraphs that can guide the development of folksonomy visualizations. We use these subgraphs to classify existing graph visualizations of folksonomies from the literature and web. We found that most works display the interrelated set of tags, while the relations between resources and/or users are largely neglected by current visualization approaches. Based on these and further findings, we discuss research challenges and potentials for future graph visualizations of folksonomies.

Categories and Subject Descriptors

H.1.0 [**Models and Principles**]: General

General Terms

Design, Standardization, Theory

Keywords

Tagging, folksonomy, representation, visualization, hypergraph, reference model, classification

1. INTRODUCTION

Tagging, i.e. the association of text labels with digital resources by users, has become a widely used indexing method in software systems. It is especially popular in social media applications, as it creates metadata and links that can be fruitfully used to classify, navigate, recommend, or retrieve digital contents. Due to this 'social' aspect, tagging in these applications is often called *social* or *collaborative tagging*. The resulting link structure has come to be known as *folksonomy* [19].

Tagging and folksonomies have also become popular research topics in the last few years. They have been analyzed and utilized in a number of works, resulting in several interesting findings, for example, on tag use and distribution [5, 6, 14]. However, apart from research on *tag clouds* [10], little systematic work has been conducted in the area of folksonomy visualization. Though there are several examples of graph visualizations for folksonomies [4, 5, 20, 17, 13, 16, 15, 9], a systematization of research in this area is missing.

We aim to make a step in this direction by presenting a reference model for the interrelated topics of representing and visualizing folksonomies as graphs. The paper shall contribute to a better understanding of the core structure resulting from tagging and support the classification and development of graph visualizations for folksonomies.

AVI'12, May 21-25, 2012, Capri Island, Italy
Copyright 2012 ACM 978-1-4503-1287-5/12/05 ...$10.00.

2. REPRESENTING TAGGING AND FOLKSONOMIES

Tagging consists of three core elements that form the basis for its representation: *resources*, *tags*, and *users*. Though these elements are differently named by different authors, their semantics and interrelation is always the same: One or more *users* (or people, actors, etc.) annotate *resources* (or objects, instances, etc.) with one or more *tags* (or keywords, labels, etc.). This fundamental principle of tagging can be defined as an axiom, at best by focusing on the resulting annotation structure:

AXIOM 1. *Each annotation resulting from tagging links exactly one resource with one user account and one or more tags. The sum of these annotations is the folksonomy.*

We can thus define three finite and disjoint sets $R = \{r_1, r_2, ..., r_k\}$, $T = \{t_1, t_2, ..., t_l\}$, and $U = \{u_1, u_2, ..., u_m\}$ that represent the resources, tags, and users. They are interconnected by the set of annotations $A = \{a_1, a_2, ..., a_n\}$ according to Axiom 1.

2.1 Three-Layer Model of Tagging

Many popular illustrations of tagging arrange the three element sets in layers that are connected by edges indicating the annotations (compare, for example, the illustrations in [5, 18, 12]). A simplified version of this three-layer model is given in Fig. 1a. It is basically an undirected tripartite graph where each partition represents one set of elements. An annotation consists of at least two binary relations in this model that connect one resource with one user via one or more tags, as given by Axiom 1. The set of annotations can thus be formally defined as $A_{2x2} \subseteq (R \times T) \cup (T \times U)$.

Although the three-layer model is an intuitive illustration of tagging, it is not suitable to represent folksonomies, as it misses information on which user annotated which resources. For instance, it is impossible to say from Fig. 1a if resource r_3 has been annotated by user u_2 or user u_3 (or even by both). This is vital information that needs to be included in any folksonomy representation [19].

To appear in:
Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI 2012)
New York, NY, USA: ACM, 2012

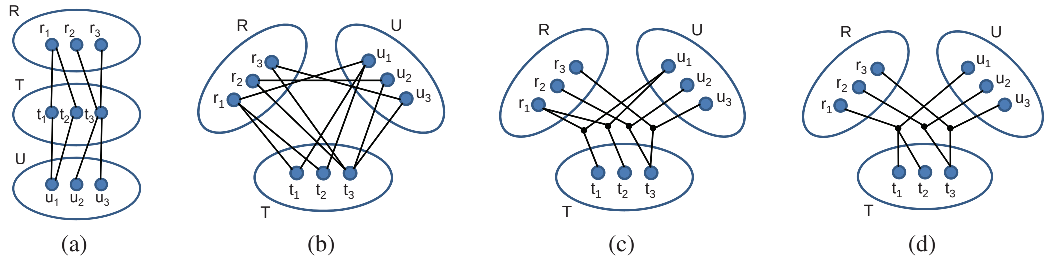

Figure 1: Formal graph representations of tagging and folksonomies, a) according to the three-layer model, and with b) relations between users and resources, c) ternary relations connecting all three element sets, d) n-ary relations connecting more than one tag.

2.2 3-Uniform Folksonomy Hypergraph

We thus need to add connections between the resource and user layers, as illustrated in Fig. 1b. The formal definition for the set of annotations would then be extended to $A_{3x2} \subseteq (R \times T) \cup (T \times U) \cup (U \times R)$. However, with this representation, it is still not possible to unambiguously express which tags have been associated by which user. In order to solve this, we can replace each set of three binary relations by one ternary relation, resulting in a 3-uniform hypergraph as illustrated in Fig. 1c. The set of all annotations can now be defined as $A_{1x3} \subseteq R \times T \times U$.

This hypergraph representation for folksonomies was first proposed by Mika [14]. Strictly speaking, it consists not only of the set of annotations A_{1x3}, i.e. the ternary relations between the sets of users, tags, and resources, but includes also the three element sets themselves. The folksonomy hypergraph can thus be defined as $G(F) = (V,E)$ with $V = R \cup T \cup U$ and $E = \{\{r,t,u\} \mid (r,t,u) \in A_{1x3}\}$. Alternatively, it can be represented in tuple form $F = (R,T,U,A_{1x3})$, as proposed by Hotho et al. [6].

Given the triadic character of this definition, various subsets can be derived by restricting the folksonomy hypergraph or quadruple to a specific element or subset of elements. For instance, the personal folksonomy (or *personomy*) P_u of a user $u \in U$ is the restriction of F to u. In tuple form, it can be defined as $P_u = (R_u, T_u, A_u)$ where $A_u = \{(r,t) \mid (r,t,u) \in A\}$ is the user's set of annotations, $T_u = \{(t) \mid (r,t) \in A_u\}$ is the user's set of tags, and $R_u = \{(r) \mid (r,t) \in A_u\}$ is the set of annotated resources.

2.3 Non-Uniform Folksonomy Hypergraph

Although the 3-uniform hypergraph and related quadruple representation provide straightforward means to represent folksonomies, they are conceptually imprecise, as the annotations are exclusively expressed by ternary relations. Annotations that consist of more than one tag must hence be split into several relations. A conceptually more accurate representation is a non-uniform hypergraph that allows for *n*-ary relations of variable *n*, as shown in Fig. 1d.

Ultimately, both hypergraph representations, the 3-uniform (Fig. 1c) and the non-uniform (Fig. 1d), contain the same information and can be transformed into each other. This is possible due to Axiom 1, stating that a resource can only be annotated once by each user. So if we know the user and resource, we can always derive the associated tags, also in the 3-uniform hypergraph representation defined above.

2.4 Beyond Tripartite Representations

There is more than the interlinked resources, tags, and users that must be considered in a comprehensive representation of tagging and folksonomies. Several ontologies have been developed in the past few years with the aim to represent further aspects of tagging (e.g. time, source, privacy, etc.). In [8], we reviewed nine available tagging ontologies, merged and classified the found concepts, and discussed modeling alternatives. Though each ontology focuses on a different aspect of tagging, they all have a similar core structure, consisting of a central class that represents either the ternary or the n-ary relations discussed in Sec. 2.2 and Sec. 2.3.

We will continue to focus on this core structure, as it provides a good starting point for the creation of graph visualizations. Other aspects of tagging are better presented differently. For instance, the time aspect can most naturally be displayed in timeline visualizations [3], while tag clouds are popular for displaying tags with their relative frequency of use.

3. DECOMPOSING THE FOLKSONOMY HYPERGRAPH

There is little point in drawing the hypergraph directly on the screen in most cases. Apart from the usual challenges of hypergraph drawing [11], the resulting visualizations would be hard to read and difficult to understand for users. We thus need to reduce the hypergraph to simpler graph structures that are limited to relevant parts of the folksonomy. Depending on what we are interested in, different subgraphs can be derived from the folksonomy hypergraph. We call this process the *decomposition of the folksonomy hypergraph*.

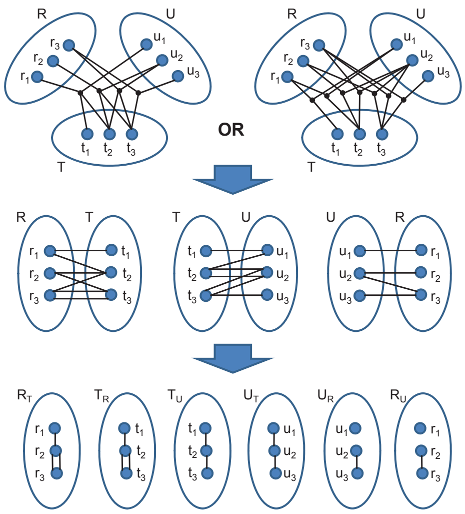

Figure 2: A complete decomposition of the folksonomy hypergraph (which is identical for the non-uniform and 3-uniform hypergraph due to Axiom 1).

Table 1: Classification of existing graph visualizations for folksonomies.

Author(s)	Subgraph	Source	Graph layout	Tag frequencies	Software (library)
Eaton [4]	$G(T_R)$	Delicious	Random force	Not shown	Java (TouchGraph)
Halpin et al. [5]	$G(T_R)$	Delicious	Random & radial force	Node size	Pajek
Zitvogel [20]	$G(T_R)_u$	Delicious	Grid & radial force	Node size	Flash
Stefaner [17]	$G(T_R)_u$	Delicious	Elastic map	Font size	Flash
Michlmayr et al. [13]	$G(T_R)_u$	Unknown	Random force	Not shown	Java (JUNG)
Shaw [16]	$G(T_U)$	Delicious	Semidef. embedding	Node & font size	Flash
Oliphant [15]	$G(RT)$, $G(RT)_u$	Flickr	Random force	Not shown	Flash (SpringGraph)
Lohmann et al. [9]	$G(RT)$	Sample data	Random force	Chain length	Flash (SpringGraph)

3.1 Complete Decomposition

Fig. 2 shows a complete decomposition of the non-uniform and 3-uniform hypergraphs discussed in Sec. 2.2 and Sec. 2.3. Compared to the examples of Fig. 1, the hypergraphs have some more edges this time to better illustrate the decomposition. Both hypergraphs contain the same information again due to Axiom 1. They are first split into three bipartite graphs that can then each be further split into a pair of unipartite graphs.

Note that each of the resulting graph pairs has the same structure, i.e. the edges in graphs $G(R_T)$ and $G(T_R)$, $G(T_U)$ and $G(U_T)$, as well as $G(U_R)$ and $G(R_U)$ are connected to the same vertices. More interestingly, each element set is represented by two graphs that represent different link types: While $G(R_T)$ links the resources based on the associated tags, $G(R_U)$ links the resources by the users who annotated them. Likewise, $G(T_R)$ and $G(T_U)$ provide different perspectives on the tag set, while $G(U_T)$ and $G(U_R)$ describe different types of links in the set of users.

Also note that all subgraphs are multigraphs, except from $G(UR)$, which is always a simple graph due to Axiom 1. This shows again that the non-uniform graph is conceptually more accurate, as it cannot be reduced to a $G(UR)$ graph with multi-edges, while this is only prevented by Axiom 1 for the 3-uniform graph.

3.2 Further Reductions of the Subgraphs

The folksonomy subgraphs can be further simplified for the visualization. A common reduction would be the transformation into weighted graphs by merging the multi-edges. In this case, the edge weights represent the number of links that exist between the vertices (e.g. the number of tag co-occurrences). Usually, such a transformation significantly reduces the number of edges, leading to a more compact and readable graph visualization.

The weighted graphs can again be further reduced. For instance, we can specify a threshold and hide all edges in the visualization that have a weight below this threshold. Alternatively, we can rank the edges by weight and visualize exclusively the top k edges. This is similar to idea of tag clouds where only a certain number of most frequently used tags is shown. Finally, further reductions can be performed by restricting the subgraphs to specific elements, as described for the folksonomy quadruple in Sec. 2.2.

4. VISUALIZING THE FOLKSONOMY SUBGRAPHS

The derived taxonomy of subgraphs can aid the development of folksonomy visualizations. It can also be helpful in classifying existing work in this area. Further categories are the folksonomy source and the graph layout, as generated by the drawing algorithm. Another more specific category is the way in which tag frequencies are visualized. Apparently, this category can only be applied to those subgraphs that contain the set of tags (i.e. graphs $G(RT)$, $G(TU)$, $G(T_R)$, and $G(T_U)$).

4.1 Classification of Existing Visualizations

In the following, we use these four categories to classify and discuss existing graph visualizations of folksonomies we found in the literature and on the web. Table 1 summarizes the results of the classification, ordered by the type of subgraph that is visualized.

4.1.1 Visualizations of Subgraph $G(T_R)$

Visualizations of the $G(T_R)$ graph are provided by Eaton [4] and Halpin et al. [5]. Both display folksonomy data from *Delicious* [1] and draw the graphs in a force-directed layout. Eaton visualizes a simple graph and provides several parameters that enable the user to control its creation. Halpin et al. adapt the node size to the tag frequency and the distance of connected nodes to the number of co-occurrences. Since both graph visualizations show all relations between the displayed tag nodes, they are very dense, with many crossing edges, and therefore hard to read.

Halpin et al. provide a second visualization of the $G(T_R)$ graph for this reason that shows only the relations of one selected tag in a radial layout. However, the visualizations of Halpin et al. are meant for illustrative purposes and do not provide any interaction.

4.1.2 Visualizations of Subgraph $G(T_R)_u$

This is different in the interactive approaches of Zitvogel [20] and Stefaner [17]. Initially, their visualizations show only the tag nodes, with the node or font size logarithmically scaled to the tag frequency, like in tag clouds. The relations of a tag with other tags are shown on demand if the user selects this tag. On the one hand, this reduces the number of visible edges and hence visual clutter, but on the other hand, it does not allow for analytical tasks which require to view the relations between all displayed tag nodes at a time, such as the visual identification of highly connected nodes.

In contrast to the previous two works, Stefaner and Zitvogel do not use common force-directed layouts to draw the graphs. Instead, Stefaner uses *elastic maps* and font weighting, resulting in a visualization that looks roughly like the inverse of a *circular tag cloud* [10]. Zitvogel uses a random grid layout that is similar to a *sequential tag cloud* [10]. If the user selects a tag, the grid changes to a radial layout that is comparable to the second visualization of Halpin et al. (see Sec. 4.1.1). Stefaner and Zitvogel use folksonomy data from Delicious again, though only from single user accounts, i.e. they visualize the restricted subgraph $G(T_R)_u$.

A third visualization of subgraph $G(T_R)_u$ is provided by Michlmayr et al. [13]. Like in the visualizations of Halpin et al., they use a force-directed layout where the node distance is roughly in inverse proportion to the number of tag co-occurrences, though they do not adapt the node weight to the tag frequency. In contrast to the other approaches, they rank the tag relations and display only the top-ranked ones and their adjacent nodes to reduce the size and density of the graph, as described in Sec. 3.2.

4.1.3 Visualization of Subgraph $G(T_U)$

We could only find a single visualization of the $G(T_U)$ graph, created by Shaw [16] from Delicious data. He uses *semidefinite embedding* and weighted nodes to draw the graph, resulting in a statistically grouped visualization that has some similarity to *clustered tag clouds* [10] enriched with tag relations. Unfortunately, the tag nodes overlap in the visualization, so that some of them are not readable.

4.1.4 Visualizations of Subgraphs G(TR) & $G(TR)_u$

Finally, there are visualizations of the $G(TR)$ graph by Oliphant [15] and Lohmann et al. [9]. Both use a force-directed layout to visualize tag relations within a set of images. While Oliphant's visualization accesses *Flickr* [2] to generate the graph, Lohmann et al. use sample data to illustrate their approach. They propose a novel type of graph representation that splits the tag nodes in the $G(TR)$ graph into chains of binary relations, which reduces the density of the graph and the number of crossing edges in most cases. However, it also increases the size of the graph and demands a transitive interpretation of the tag relations.

5. DISCUSSION

As we can see from the classification, most of the existing work visualizes either the graph $G(T_R)$ graph or its restricted subgraph $G(T_R)_u$ (cp. Table 1). These 'tag graph' visualizations are related to tag clouds as they show the set of tags, but they additionally display relations between the tags based on their co-use. There are also visualizations of the $G(T_U)$, $G(RT)$ and $G(RT)_u$ graphs, but we could not find any visualization of the other folksonomy subgraphs derived in Sec. 3.

However, visualizations for these subgraphs can provide valuable insight. For instance, a visual analysis of the $G(TU)$ and $G(UR)$ graphs may uncover interesting links between users based on the tags they use or the resources they annotate. In social media contexts, visualizations of the $G(UR)$ and $G(R_U)$ graphs may be good alternatives to the $G(RT)$ and $G(R_T)$ graphs, as they do not show tag-based but 'social' relations between digital resources. Though these subgraphs are already used in folksonomy-based filtering and recommendation [7], they have not yet been examined in folksonomy visualizations to the best of our knowledge. It would also be interesting to see visualizations for other folksonomy sources than Delicious, which is used by most of the reviewed approaches (cp. Table 1).

Further research challenges are the size and density of folksonomy graphs that must be significantly reduced for readable visualizations. We presented different strategies, from visualizing only selected nodes and edges (e.g. the top weighted ones) to letting the user adapt the size and density of the graph. A related research question is the optimal mapping of the subgraph structures to a two-dimensional plane. While most of the reviewed works use simple force-directed layouts, more advanced drawing algorithms are likely to provide better results.

6. ACKNOWLEDGMENTS

This work was partially funded by the Spanish Ministry of Science and Innovation within the UrThey project (TIN2009-09687).

7. REFERENCES

[1] Delicious. http://delicious.com.

[2] Flickr. http://www.flickr.com.

[3] M. Dubinko, R. Kumar, J. Magnani, J. Novak, P. Raghavan, and A. Tomkins. Visualizing tags over time. In *Proceedings of the 15th International Conference on World Wide Web*, WWW '06, pages 193–202. ACM, 2006.

[4] A. Eaton. Graph del.icio.us related tags. http://hublog.hubmed.org/archives/001049.html, 2005.

[5] H. Halpin, V. Robu, and H. Shepherd. The complex dynamics of collaborative tagging. In *Proceedings of the 16th International Conference on World Wide Web*, WWW '07, pages 211–220. ACM, 2007.

[6] A. Hotho, R. Jäschke, C. Schmitz, and G. Stumme. Information retrieval in folksonomies: Search and ranking. In *Proceedings of the 3rd European Semantic Web Conference*, ESWC '06, pages 411–426. Springer, 2006.

[7] R. Jäschke, L. B. Marinho, A. Hotho, L. Schmidt-Thieme, and G. Stumme. Tag recommendations in folksonomies. In *Proceedings of the 11th European Conference on Principles and Practice of Knowledge Discovery in Databases*, PKDD '07, pages 506–514. Springer, 2007.

[8] S. Lohmann, P. Díaz, and I. Aedo. MUTO: The modular unified tagging ontology. In *Proceedings of the 7th International Conference on Semantic Systems*, I-SEMANTICS '11, pages 95–104. ACM, 2011.

[9] S. Lohmann, P. Heim, L. Tetzlaff, T. Ertl, and J. Ziegler. Exploring relationships between annotated images with the chaingraph visualization. In *Proceedings of the 4th International Conference on Semantic and Digital Media Technologies*, SAMT '09, pages 16–27. Springer, 2009.

[10] S. Lohmann, J. Ziegler, and L. Tetzlaff. Comparison of tag cloud layouts: Task-related performance and visual exploration. In *Proceedings of the 12th International Conference on Human-Computer Interaction*, INTERACT '09, pages 392–404. Springer, 2009.

[11] E. Mäkinen. How to draw a hypergraph. *International Journal of Computer Mathematics*, 34(3):177–185, 1990.

[12] C. Marlow, M. Naaman, D. Boyd, and M. Davis. Ht06, tagging paper, taxonomy, flickr, academic article, to read. In *Proceedings of the 17th ACM Conference on Hypertext and Hypermedia*, HYPERTEXT '06, pages 31–40. ACM, 2006.

[13] E. Michlmayr and S. Cayzer. Learning user profiles from tagging data and leveraging them for personal(ized) information access. In *Proceedings of the WWW 2007 Workshop on Tagging and Metadata for Social Information Organization*, 2007.

[14] P. Mika. Ontologies are us: A unified model of social networks and semantics. In *Proceedings of the 4th International Semantic Web Conference*, ICWE '05, pages 522–536. Springer, 2005.

[15] F. Oliphant. TagGraph. http://taggraph.com, 2005.

[16] B. Shaw. Semidefinite embedding applied to visualizing folksonomies. http://www.metablake.com/advml/adv-ml-project.pdf, 2005.

[17] M. Stefaner. Visual tools for the socio-semantic web. Master thesis, University of Applied Sciences Potsdam, 2007.

[18] T. Vander Wal. Explaining and showing broad and narrow folksonomies. http://www.personalinfocloud.com/2005/02/explaining_and_.html, 2005.

[19] T. Vander Wal. Folksonomy coinage and definition. http://vanderwal.net/folksonomy.html, 2007.

[20] O. Zitvogel. Delicious Soup. www.zitvogel.com/delicioussoup, 2005.