

EMERGENT COMMUNITY STRUCTURE IN SOCIAL TAGGING SYSTEMS

CIRO CATTUTO^{*,†,‡,§}, ANDREA BALDASSARRI[†], VITO D. P. SERVEDIO^{*,†} and VITTORIO LORETO^{†,‡}

*Museo Storico della Fisica e Centro Studi e Ricerche "Enrico Fermi", Compendio Viminale, 00184 Roma, Italy

[†]Dipartimento di Fisica, Università di Roma "La Sapienza", P.le A. Moro, 2, 00185 Roma, Italy

[‡]Institute for Scientific Interchange (ISI), Torino, Italy [§]ciro.cattuto@roma1.infn.it

> Received 27 September 2007 Revised 15 March 2008

A distributed classification paradigm known as *collaborative tagging* has been widely adopted in new Web applications designed to manage and share online resources. Users of these applications organize resources (Web pages, digital photographs, academic papers) by associating with them freely chosen text labels, or *tags*. Here we leverage the social aspects of collaborative tagging and introduce a notion of *resource distance* based on the collective tagging activity of users. We collect data from a popular system and perform experiments showing that our definition of distance can be used to build a weighted network of resources with a detectable community structure. We show that this community structure clearly exposes the semantic relations among resources. The communities of resources that we observe are a genuinely emergent feature, resulting from the uncoordinated activity of a large number of users, and their detection paves the way for mapping emergent semantics in social tagging systems.

 $Keywords\colon$ Folksonomy; collaborative tagging; emergent semantics; online communities; Web 2.0.

1. Introduction

Information systems on the World Wide Web have been increasing in size and complexity to the point where they presently exhibit features typically attributed to *bona fide* complex systems. They display rich, high-level behaviors that are causally connected in nontrivial ways to the dynamics of their interacting elementary parts. Because of this, concepts and formal tools from the science of complex systems can play an important role in understanding the structure and dynamics of such systems.

REST web services Good intro to the REST "architecture" to web service tutorial guidelines api rest by hotho and 3 other people on 2006-04-04 16:11:47 copy

Fig. 1. The basic unit of information in a folksonomy, i.e. a post, is shown as it appears in the interface of bibsonomy.org, a social collaborative tagging system for bookmarks and scientific references. At the top, the title of the resource (a Web page) is shown, followed by its own subtitle. Then the list of tags associated by the user hotho is displayed. Other pieces of information are: the number of other users who inserted the same resource in the system, and the date and time of insertion of the present post.

This study focuses on the recently established paradigm of collaborative tagging [1, 2]. In Web applications like del.icio.us,^a Flickr^b and BibSonomy,^c users organize diverse *resources* — ranging from Web pages to academic papers and photographs — with semantically meaningful information in the form of text labels, or *tags*. Tags are freely chosen and users associate resources with them in a totally uncoordinated fashion. Nevertheless, the tagging activity of each user is globally visible to the user community and the tagging process develops genuine social aspects and complex interactions [3,4], eventually leading to a bottom-up categorization of resources shared throughout the user community. The open-ended set of tags used within the system — commonly referred to as "folksonomy" — can be used as a sort of semantic map to navigate the contents of the system itself.

In Fig. 1 a single annotation example (called a "post") is shown as it appears in the interface of the bibsonomy.org system.

Our work is based on experimental data from one of the largest and most popular collaborative tagging systems, del.icio.us, currently used by over a million users to manage and share their collections of Web bookmarks.

The main point of our work is neither to present a new spectral community detection algorithm, nor to report a large data set analysis. Rather, we want to show that, choosing the right projection and the right weighting procedure, we can produce a weighted, undirected network of resources from the full tripartite folksonomy network, which embeds a meaningful social classification of resources. This is especially surprising, considering that users annotate resources in a very anarchic, uncoordinated and noisy way.

In Sec. 2 we describe the experimental data we collected. In Sec. 3 we introduce a notion of resource distance based on the collective activity of users. Based on that, we set up an experiment using actual data from del.icio.us and we build a weighted network of resources. In Sec. 4 we show that spectral methods from complex network theory can be used to detect clusters of resources in the above network, and we

^ahttp://del.icio.us ^bhttp://flickr.com ^chttp://www.bibsonomy.org characterize those clusters in terms of user tags, exposing semantics. Finally, Sec. 5 gives an overview of our results and points to directions for future work.

2. Experimental Data

Our analysis focuses on del.icio.us for several reasons: (i) it was the first system to deploy the ideas of collaborative tagging on a large scale, so it has acquired a paradigmatic character is the natural starting point for any quantitative study; (ii) it has a large user community and contains a huge amount of raw data on the structure and dynamics of a folksonomy; (iii) it is a broad folksonomy [5], i.e. single tag associations by different users retain their identity and can be individually retrieved. This allows us to measure the number of times that a given tag X was associated with a specific resource as the number (f_X) of users who established that resource-tag association (see also Fig. 2). That is to say, a broad folksonomy has a natural notion of weight for tag associations, which is based on social agreement. On studying delicious we adopt a resource-centric view of the system, i.e. we investigate the emergent correspondence between a given resource and the tags that all users associate with it. We factor out the detailed identity of the users and deal only with the set of tags associated by the user community with a given resource, as well as with the frequencies of occurrence of those tags in the context of the resource.



Fig. 2. The collective activity of users associates with each resource a weighted set of tags, where the weight of a tag is given by its frequency of occurrence in the context of a resource. The weighted set of tags is commonly visualized by using a graphical device called a *tag cloud*: the most frequent tags associated with a given resource are shown, and the font size of each tag is proportional to the logarithm of its frequency of occurrence. Our definition of similarity w_{R_1,R_2} [Eq. (1)] measures the weighted overlap between the tag clouds associated with the resources R_1 and R_2 . Tags marked in red belong to $T_1 \cap T_2$, the set of tags shared by the two resources.

To collect data, we used a Web crawler that connects to del.icio.us and navigates the system's interface as an ordinary user would do, extracting tagging metadata and storing it for further postprocessing. Our client connects to del.icio.us and downloads the Web pages associated with a given set of resources, using an HTML parser to extract the tagging information from the page. The system allows one to get the complete set of annotations associated with each resource. The data used for the present analysis were retrieved in October 2006.

3. Resource Networks from Collective Tagging Patterns

In a collaborative tagging system, a set of resources defines a "semantic space" that is explored and mapped by a community of users, as they bookmark and tag those resources [6]. We want to investigate whether the tagging activity is actually structuring the space of resources in a semantically meaningful way, i.e. whether partitions or subsets of resources emerge, associated with tagging patterns that point to well-defined meanings, areas of interest or topics. These groups of resources could also identify, in principle, communities of users sharing the same view of resources, or the same emergent vocabulary.

In order to gain insight into the above problem, we set up an experiment using del.icio.us as a data source. We want to stress here that, since the aim of the work is to investigate whether an emergent community structure exists in folksonomy data, we are not concerned with the completeness of the dataset used. Rather, we decided to perform the experiment on the following subset: we selected two popular tags that appear to be semantically unrelated (design and politics), and for each of them we extracted from del.icio.us a set of 200 randomly chosen resources (we took the first 200 returned by the system, representing the most recently introduced by users). For each resource, we collected the complete set of annotations, i.e. all the tag assignments relative to that resource. The corresponding dataset used for this experiment thus consists of 400 resources: half of them have been associated with the tag design, while the other half have been tagged with *politics*. The idea is to construct a dataset containing at least two semantically well-separated subsets. For each resource in the dataset, the entire tagging history was retrieved from del.icio.us, so that all the tag associations involving the chosen 400 resources are known. In other words, we know how the entire user community of del.icio.us "categorized" the selected resources in terms of freely chosen tags, with no biases due to data collection.

To uncover structures linked to specific tagging patterns, we introduce a notion of similarity between resources based on how those resources were tagged by the user community. For each resource, we define a *tag cloud* as the weighted set of tags that have been used to bookmark that resource, where the weight of tag t is its frequency of occurrence f_t in the context of that resource (Fig. 2). We want to formalize the intuitive idea that two resources are similar if the corresponding tag clouds have a high degree of overlap. Given two generic resources R_1 and R_2 , and the corresponding sets of tags T_1 and T_2 , a natural measure of tag cloud overlap would be the standard set overlap given by the cardinality of the intersection set $T_1 \cap T_2$ divided by the cardinality of the union set $T_1 \cup T_2$. This simple measure, however, has a major fault: since no notion of tag weight (frequency) is used, it is not sensitive to the social aspects of tagging encoded in tag frequencies (and, as such, it is also vulnerable to tagging noise, i.e. errant, strange, incorrect or even malicious tagging, or spamming [12–14]).

To overcome this limitation we adopt a term frequency–inverse document frequency (TF–IDF) weighting procedure [7]. The TF–IDF weight is commonly used in information retrieval and text mining, and represents a statistical measure used to evaluate how specific a term is in identifying a document belonging to a collection of documents. The importance of a term increases proportionally to the number of times the term appears in the document, and inversely proportionally to the global frequency of the same term in the document collection.

We denote with f_t^1 and f_t^2 the frequencies of occurrence of tag t in T_1 and T_2 , respectively, and with f_t the global frequency of tag t, i.e. the total number of times that tag t was used in association with all the resources under study.

In the spirit of the TF–IDF techniques, we normalize the frequencies of tags by their global frequencies. When a tag is shared by the resources R_1 and R_2 , it has two different frequencies, f_t^1 in the context of R_1 and f_t^2 in the context of R_2 . When performing the intersection between tag clouds, we use the lowest of those frequencies to define the weight of tag t in the intersection set $T_1 \cap T_2$, while we use the highest of those frequencies when weighting the contribution of the same tag in the union set $T_1 \cup T_2$. More precisely, we define the similarity between R_1 and R_2 as

$$w_{R_1,R_2} = \frac{\sum_{t \in T_1 \cap T_2} \frac{\min(f_t^1, f_t^2)}{f_t}}{\sum_{t \in T_1 \cap T_2} \frac{\max(f_t^1, f_t^2)}{f_t} + \sum_{t \in T_1 - T_2} \frac{f_t^1}{f_t} + \sum_{t \in T_2 - T_1} \frac{f_t^2}{f_t}}{f_t}.$$
 (1)

The above expression is an extension of the simple measure of set overlap, where the numerator is a weighted form of set intersection and the denominator is a weighted form of set union. By definition, $0 \le w_{R_1,R_2} \le 1$. Of course, the above definition is just one of the possible similarity measures that can be employed, and the validation of the measure we introduce here is left to the results obtained by using it, as shown in Sec. 4. The similarity matrix introduced above can be regarded as the adjacency matrix of a weighted network of resources [8], where w_{R_1,R_2} is the strength of the edge connecting nodes R_1 and R_2 .

Figure 3 shows the distribution of similarities (edge strengths in the weighted network) among all the pairs of resources, for three different sets of resources: the subset of resources sharing the tag *design*, the subset of resources sharing the tag *politics* and the union of those two subsets. Notice that the global frequency f_t of a given tag t depends on the set of resources chosen for the analysis. From the plot it



Fig. 3. Probability distributions of link strengths. The logarithmically binned histogram of link strengths for all pairs of resources within a given set is displayed for three sets of resources: empty squares correspond to resources tagged with *design*, filled squares correspond to resources tagged with *politics*, and blue circles correspond to the union of those two sets. It is important to observe that strength values span several orders of magnitude, so that a nonlinear function of link strengths becomes necessary in order to capture the full dynamic range of strength values.

is evident that weights span a wide range of values and the logarithm of the weight is best-suited to appreciating the full range of strength values.

4. Community Structure of the Resource Network

In order to investigate the existence of underlying structures in the set of resources, we proceed as follows. First, we transform the similarity matrix w_{R_1,R_2} in order to compress the dynamic range of strength values. Since the logarithmic scale gives a good representation of the strength variability (Fig. 3), but has divergence problems in the neighborhood of zero, we consider a matrix where each element is raised to a small (arbitrary) power, $\gamma = 0.1$. Thus, the similarity matrix w' which we will use in the following is defined as

$$w'_{R_1,R_2} = (w_{R_1,R_2})^{\gamma}.$$
(2)

Note that the similarity metric 2 is similar to that introduced in Refs. 15 and 16 for a clustering experiment in an ontology of Web pages, and was inspired by information theory arguments.

Figure 4 displays the similarity matrix (link strengths of the weighted similarity network) between pairs of resources w'_{R_1,R_2} for the full set of 400 resources. The resources are randomly ordered and no structures are visible in this representation.



Fig. 4. Matrix w' of link strengths [Eq. (2)] for the entire set of 400 randomly ordered resources. Except for the bright diagonal, whose elements are identically equal to 1 because of the normalization property of the strength w, the matrix appears featureless. Note that no community structure appears. Color figure is only available in electronic version.

The problem we have to tackle now is finding the sequence of row and column permutations of the similarity matrix that permits one to visually identify the presence of communities of resources, if at all possible. The goal is to obtain a matrix with a clear visible block structure on its main diagonal. One possible way to approach this problem is to construct an auxiliary matrix and use information deduced from its spectral properties to rearrange rows and columns of the original matrix. The quantity we consider is the matrix

$$Q = S - W, (3)$$

where $W_{ij} = (1 - \delta_{ij})w'_{ij}$ and S is a diagonal matrix where each element on the main diagonal equals the sum of the corresponding row of W, i.e. $S_{ij} = \delta_{ij} \sum_{j} W_{ij}$. The matrix Q is nonnegative and resembles the Laplacian matrix of graph theory. As shown in Refs. 9 and 10, the study of its spectral properties can reveal the community structure of the network.

The main idea is to consider the lowest eigenvalues of Q. According to the definition of Q, there is a always a zero eigenvalue corresponding to an eigenvector with equal components, i.e. a trivial constant eigenvector. Let us now consider the simple case where the matrix Q is composed of exactly two nonzero blocks along its main diagonal (i.e. with two clearly separated semantic communities). In this case, two eigenvectors with zero eigenvalue are present, signalling the existence of two disconnected components. When nonzero entries connecting the two blocks are present, only one null eigenvalue survives, and the components of the eigenvectors with the lowest eigenvalues reveal the community structure. Given the set of these nontrivial eigenvectors, a very simple way to identify the communities consists in plotting their components on a (multidimensional) scatter plot. Each axis reports the values of the components of the eigenvectors. In particular, each point has coordinates equal to the homologous components of one eigenvector. In this kind of plot, communities emerge as well-defined clusters of points aligned along specific directions. The components involved in each cluster identify the elements belonging to a given community. Once the communities are identified, it is interesting to permute the indices of the original matrix W such that the components of the same community become adjacent. The corresponding matrix should appear roughly madeup of diagonal blocks, possibly with mixing terms signalling an overlap between communities (blocks).

Figure 5 displays the eigenvalues of Q sorted by their value. As expected, the null eigenvalue is present, corresponding to the trivial constant eigenvector.

Figure 6 displays a three-dimensional scatter plot illustrating the structure of the three eigenvectors that correspond to the three lowest nontrivial eigenvalues of Q (the second, third and fourth ones; see Fig. 5).

The axes show the values of the components of the second, third and fourth eigenvectors, respectively (denoted by V_2 , V_3 and V_4). In particular, each point has coordinates equal to the homologous components for the three nontrivial eigenvectors considered.

The existence of at least five well-defined communities is evident, with each community corresponding to one of the five well-separated nonnull eigenvalues of Fig. 5.



Fig. 5. Eigenvalues of the matrix Q [Eq. (3)]. Resource communities correspond to nontrivial eigenvalues of the spectrum, such as the ones visible on the leftmost side of the plot and in the inset. The three eigenvalues marked in the inset correspond to the eigenvectors plotted in Fig. 6.



Fig. 6. Eigenvectors of the matrix Q [Eq. (3)]. The scatter plot displays the component values of the first three nontrivial eigenvectors of the matrix (marked with circles in Fig. 5). The scatter plot is parametric in the component index. Five or six clusters are visible, corresponding to the smallest nontrivial eigenvalues of the similarity matrix. Each cluster, marked with a numeric label, defines a community of "similar" resources (in terms of tag clouds). Blue and red points correspond to resources tagged with *design* and *politics*, respectively. Notice that our approach clearly recovers the two original sets of resources, and also highlights a few finer-grained structures. Tag clouds for the identified communities are shown in Fig. 8.

A sixth, very small community, corresponding to the sixth nontrivial eigenvalue, is barely visible.

Once we have diagonalized the matrix Q the permutation of indices necessary to sort the component values of these eigenvectors yields the desired ordering of rows and columns in the original matrix W. By performing this reordering it is possible to visualize the matrix of strengths of Fig. 4 in a way that makes it maximally diagonal. Figure 7 reports the reordered matrix.

An interesting question is now whether the communities we have found correspond to semantic differences in the set of resources. In order to check this point we build for each community a tag cloud from the tags associated with the corresponding group of resources.

Figure 8 reports the six tag clouds (ordered by decreasing number of member resources), where the font size of each tag, as usual, is proportional to the logarithm of its frequency of occurrence. Despite the intrinsic difficulty of identifying the semantic context defined by a given tag cloud, it is possible to recognize that each comunity of resources — at least for the four largest — comprises resources with a specific semantic connotation. In particular, the first community can be associated with humor in politics, the second with visual design, the third with political blogs and the fourth with Web design.



Fig. 7. Matrix w' of link strengths [see Eq. (2)] for our set of 400 resources. Here the resource indices are ordered by community membership (the sequence of communities along the axes is 2, 4, 6, 5, 3, 1; see Fig. 8). In striking contrast to Fig. 4, the permutation of indices we employed clearly exposes the community structure of the set of resources: two large groups of resources with high similarity, corresponding to the blue/red rectangles at the top right and bottom left of the matrix, correspond respectively to resources tagged with *design* and *politics*. On top of this, our approach reveals the presence of finer-grained community structures within the above communities (red rectangular regions toward the center of the matrix). On direct inspection, these communities of resources turn out to have a rather well-defined semantic characterization in terms of tags, as shown by the tag clouds of Fig. 8. Color figure is only available in electronic version.



Fig. 8. Tag clouds for the six resource communities identified by our analysis (see Fig. 6) ordered by decreasing community size. Each tag cloud shows the 30 most frequent tags associated with resources belonging to the corresponding community. As usual, the size of text labels is proportional to the logarithm of the frequency of the corresponding tag. The first two communities (the biggest ones) correspond largely to the main division between resources tagged with *politics* and *design*, respectively. Notice how each tag cloud is strongly characterized by only one of those two tags. In addition to discriminating the above two main communities, our approach identifies additional, unexpected communities. On inspecting the corresponding tag clouds, one can recognize a rather well-defined semantic connotation pertaining to each community, as discussed in the main text.

5. Conclusions

The increasing impact of Web-based social tools for the organization and sharing of resources is motivating new research at the frontier of complex systems science and computer science, with the goal of harvesting the emergent semantics [11] of these new tools.

The increasing interest in such new tools is based on the belief that the anarchic, uncoordinated activity of users can be used to extract meaningful and useful information. For instance, in social bookmarking systems, people annotate a personal list of resources with freely chosen tags. Whether or not this could provide a "social" classification of resources is the point we want to investigate with this work. In other words, we investigate whether an emergent community structure exists in folksonomy data. To this end, we focused on a popular social bookmarking system and introduced a notion of similarity between resources (annotated objects) in terms of social patterns of tagging. We used our notion of similarity to build weighted networks of resources, and showed that spectral community detection methods can be used to expose the emergent semantics of social tagging, identifying well-defined communities of resources that appear associated with distinct and meaningful tagging patterns. The present analysis was limited to an experiment where the set of resources was artificially built by selecting resources tagged with semantically unrelated tags: future directions for this research include large-scale experiments on broader sets of resources, to assess the robustness of our method, as well as the investigation of other indicators of social agreement that can be leveraged to expose structures in folksonomies. Such efforts could lead to improved user interfaces, increasing both the usability and utility of these new, powerful tools.

Acknowledgments

The authors wish to thank Melanie Aurnhammer, Andreas Hotho and Gerd Stumme for very interesting discussions. This research has been partly supported by the TAGora project funded by the Future and Emerging Technologies program (IST-FET) of the European Commission under the contract IST-34721. The information provided is the sole responsibility of the authors and does not reflect the Commission's opinion. The Commission is not responsible for any use that may be made of data appearing in this publication.

References

- Mates, A., Folksonomies: Cooperative Classification and Communication Through Shared Metadata (Computer Mediated Communication, LIS590CMC, Graduate School of Library and Information Science, University of Illinois at Urbana-Champaign, 2004).
- [2] Hammond, T., Hannay, T., Lund, B. and Scott, J., Social bookmarking tools (I): A general review, *D-Lib Mag.* 11(4) (2005).
- [3] Golder, S. and Huberman, B. A., Usage patterns of collaborative tagging systems, J. Inf. Sci. 32 (2006) 198.

- [4] Cattuto, C., Loreto, V. and Pietronero, L., Semiotic dynamics and collaborative tagging, Proc. Natl. Acad. Sci. U.S.A. 104 (2007) 1461.
- [5] Vander Wal, T., Explaining and showing broad and narrow folksonomies. http://www.personalinfocloud.com/2005/02/explaining_and_.html (2005).
- [6] Hotho, A., Jäschke, R., Schmitz, C. and Stumme, G., Emergent semantics in BibSonomy, in *Proc. Workshop on Applications of Semantic Technologies*, eds. Hochberger, C. and Liskowsky, R. (Informatik für Menschen, 2006), Band 2, p. 94.
- [7] Salton G. and McGill, M. J., Introduction to Modern Information Retrieval (McGraw-Hill, 1983).
- [8] Barrat, A., Barthelemy, M., Pastor-Satorras, R. and Vespignani, A., The architecture of complex weighted networks, *Proc. Natl. Acad. Sci. U.S.A.* 101 (2004) 3747.
- [9] Capocci, A., Servedio, V. D. P., Caldarelli, G. and Colaiori, F., Physica A 352, 669 (2005).
- [10] Newman, M. E. J., *Phys. Rev. E* **74** (2006) 036104.
- [11] Steels, L., Semiotic dynamics for embodied agents, *IEEE Intell. Syst.* **21** (2006) 32.
- [12] Cattuto, C., Schmitz, C., Baldassarri, A., Servedio, V. D. P. Loreto, V., Hotho, A., Grahl, M. and Stumme, G., Network properties of folksonomies, *AI Commun. J.*, Special Issue on "Network analysis in natural sciences and engineering," eds. Hoche, S., Nürnberger, A. and Flach, J., **20**(4) (2007) 245–262.
- [13] Koutrika, G., Effendi, F. A., Gyöngyi, Z., Heymann P. and Garcia-Molina, H., Combating spam in tagging systems, in *AIRWeb '07: Proc. 3rd Int. Workshop on Adver*sarial Information Retrieval on the Web (ACM Press, 2007), pp. 57–64.
- [14] Heymann, P., Koutrika, G. and Garcia-Molina, H., Fighting spam on social web sites: A survey of approaches and future challenges, *IEEE Internet Comput.*, **11** (2007) 35–46.
- [15] Maguitman, A. G., Menczer, F., Roinestad, H. and Vespignani, A., Algorithmic detection of semantic similarity, in WWW '05: Proc. 14th Int. Conf. World Wide Web (2005), pp. 107–116.
- [16] Maguitman, A., Menczer, F., Erdinc, F., Roinestad, H. and Vespignani, A., Algorithmic computation and approximation of semantic similarity, World Wide Web 9 (2006) 431–456.