

HT06, Tagging Paper, Taxonomy, Flickr, Academic Article, ToRead

Cameron Marlow¹, Mor Naaman¹, danah boyd^{1,2}, Marc Davis^{1,2}

¹Yahoo! Research Berkeley
1950 University Avenue, Suite 200
Berkeley, CA 94704-1024

²UC Berkeley School of Information
102 South Hall
Berkeley, CA 94720-4600

{cameronm, mor, danah, marcd}@yahoo-inc.com

ABSTRACT

In recent years, *tagging systems* have become increasingly popular. These systems enable users to add keywords (i.e., “tags”) to Internet resources (e.g., web pages, images, videos) without relying on a controlled vocabulary. Tagging systems have the potential to improve search, spam detection, reputation systems, and personal organization while introducing new modalities of social communication and opportunities for data mining. This potential is largely due to the social structure that underlies many of the current systems.

Despite the rapid expansion of applications that support tagging of resources, tagging systems are still not well studied or understood. In this paper, we provide a short description of the academic related work to date. We offer a model of tagging systems, specifically in the context of web-based systems, to help us illustrate the possible benefits of these tools. Since many such systems already exist, we provide a taxonomy of tagging systems to help inform their analysis and design, and thus enable researchers to frame and compare evidence for the sustainability of such systems. We also provide a simple taxonomy of incentives and contribution models to inform potential evaluative frameworks. While this work does not present comprehensive empirical results, we present a preliminary study of the photo-sharing and tagging system Flickr to demonstrate our model and explore some of the issues in one sample system. This analysis helps us outline and motivate possible future directions of research in tagging systems.

Categories and Subject Descriptors

H.1.1 [Information Systems]: Models and Principles – *Systems and Information Theory*.

General Terms

Algorithms, Design, Human Factors.

Keywords

Tagging systems, taxonomy, folksonomy, tagsonomy, Flickr, categorization, classification, social networks, social software, models, incentives, research.

1. INTRODUCTION

Web-based *tagging systems* such as Del.icio.us, Technorati and Flickr allow participants to annotate a particular *resource*, such as a web page, a blog post, an image, a physical location, or just about any imaginable object with a freely chosen set of keywords

(“tags”). In this paper, we aim to articulate a framework for studies of such systems.

One approach to tagging has emerged in “social bookmarking” tools where the act of tagging a resource is similar to categorizing personal bookmarks. In this model, tags allow users to store and collect resources and retrieve them using the tags applied. Similar keyword-based systems have existed in web browsers, photo repository applications, and other collection management systems for many years; however, these tools have recently increased in popularity as elements of social interaction have been introduced, connecting individual bookmarking activities to a rich network of shared tags, resources, and users.

Social tagging systems, as we refer to them, allow users to share their tags for particular resources. In addition, each tag serves as a link to additional resources tagged the same way by others. Because of their lack of predefined taxonomic structure, social tagging systems rely on shared and emergent social structures and behaviors, as well as related conceptual and linguistic structures of the user community. Based on this observation, the popular tags in social tagging systems have recently been termed *folksonomy* [22], a folk taxonomy of important and emerging concepts within the user group.

Social tagging systems may afford multiple added benefits. For instance, a shared pool of tagged resources enhances the metadata for all users, potentially distributing the workload for metadata creation amongst many contributors. These systems may offer a way to overcome the Vocabulary Problem – first articulated by George Furnas et al in [8] – where different users use different terms to describe the same things (or actions). This disagreement in vocabulary can lead to missed information or inefficient user interactions. The taxonomy of tagging systems articulated in this paper, and the results of our preliminary experiments on the relationship between tag overlap and social connection, both point to the possibility that thoughtful sociotechnical design of tagging systems may uncover ways to overcome the Vocabulary Problem without requiring either the rigidity and steep learning curve of tightly controlled vocabularies, or the computational complexity and relatively low success of purely automatic approaches to term disambiguation.

Figure 1 shows a conceptual model for social tagging systems. In this model, users assign tags to a specific resource; tags are represented as typed edges connecting users and resources. Resources may be also be connected to each other (e.g., as links between web pages) and users may be associated by a social network, or sets of affiliations (e.g., users that work for the same company).

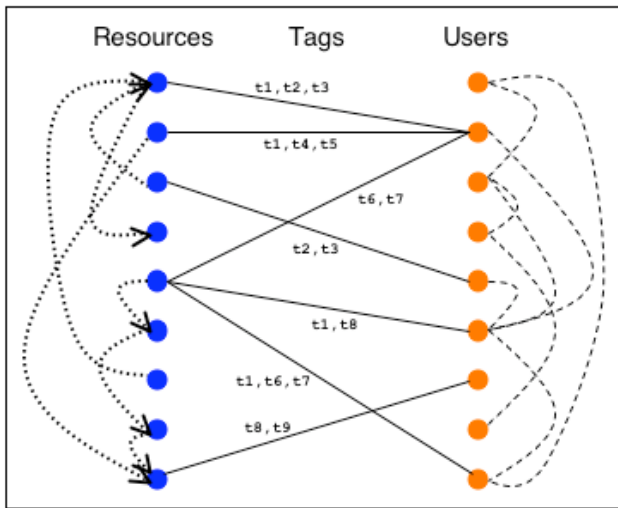


Figure 1. A model of tagging systems.

The three individual elements of the model depicted in Figure 1 have been studied independently in the past, usually in the context of web-based systems:

- *Resources.* The relationship between resources and links is a well-researched area. Most prominently, PageRank [18] has made analysis of link structure on the web a household name.
- *Users.* Analysis of social ties and social networks is an established subfield of sociology [25] and has received attention from physicists, computer scientists, economists, and numerous other areas of study.
- *Tags.* Recently, the aggregation and semantic aspects of tags have been discussed and debated at length [16]. This discussion has mainly focused on the quality of information produced by tagging systems and the possible tradeoffs between folksonomies and crafted ontologies [17, 20]. Furthermore, the challenges of shared vocabularies for description have been studied in the information science and library science communities for many years [8].

Despite these individual contributions (which we will revisit in more detail in Section 2), to fully understand tagging systems we believe a holistic approach is necessary. Walker [24] describes tagging as “feral hypertext”, a structure out of control, where the same tag is assigned to different resources with different semantic senses, and thus associates otherwise unrelated resources. However, by considering the entire model, computer systems could make inferences that “domesticate” (to use Walker’s terms) these “feral” tags. For example, tag semantics and synonyms could potentially be inferred by analyzing the structure of the social network, and identifying certain portions of the network that use certain tags for the same resource, or related resources, interchangeably. These tags may be synonymous.

A unified user-tag-resource approach might be useful for many key web technologies, including: search and information retrieval; information organization, discovery and communication; spam filtering; reducing effects of link spam, and improving on trust metrics; identifying trends and emerging topics globally and within communities; and locating experts and opinion leaders in specific domains.

Variations in the model described in Figure 1 are possible. For example, links between resources could be absent, and likewise for users. Nevertheless, in these circumstances, we can still

observe connections between users, tags, and resources. These connections define an implicit relationship between resources through the users that tag them; similarly, users are connected by the resources they tag.

In order to better frame the space of social tagging systems, we describe two organizational taxonomies for social tagging systems, developed by analyzing and comparing the design and features of many existing social tagging systems. The taxonomies describe:

- *System design and attributes.* We claim that the place of a tagging system in this taxonomy will greatly affect the nature and distribution of tags, and therefore the attributes of the information collected by the system.
- *User incentives.* User behaviors are largely dictated by the forms of contribution allowed and the personal and social motivations for adding input to the system. The place of a tagging system in this taxonomy will affect its overall characteristics and benefits.

To demonstrate how these classifications affect the properties of tags and users, we will present a study of Flickr, one of the most popular tagging systems on the web today. We compare our findings from the Flickr study to the work of Golder and Huberman [9] on Del.icio.us. Flickr and Del.icio.us are complementary examples of tagging-systems in our taxonomies; we present initial evidence that the dynamics of these systems are quite different.

To this end, the next section provides details on related work, mostly concentrating on academic research work in related areas. Section 3 briefly outlines a number of current tagging systems used as illustration in different parts of this paper. Section 4 describes our taxonomies of tagging system design choices and incentives. In Section 5 we present the results of our study of tagging in Flickr, a photo-sharing tagging system. We present a summary and outline future directions of research in Section 6.

2. RELATED WORK

Despite a considerable amount of attention in academic circles, as represented in various blog posts [17,20], little academic research work has been invested in tagging systems to date.

Perhaps the most significant formal study of tagging systems appears in the work of Golder and Huberman [9]. The authors study the information dynamics in “collaborative tagging systems”—specifically, the Del.icio.us system. The authors discuss the information dynamics in such a system, including how tags by individual users are used over time, and how tags for an individual resource (in the case of Del.icio.us, web resources) change—or more specifically, stabilize—over time. We refer to their findings again in Section 5.

Golder and Huberman also discuss the semantic difficulties of tagging systems. As they point out, *polysemy* (when a single word has multiple related meanings) and *synonymy* (when different words have the same meaning) in the tag database both hinder the precision and recall of tagging systems. In addition, the different expertise and purposes of tagging participants may result in tags that use various levels of abstraction to describe a resource: a photo can be tagged at the “basic level” of abstraction [14] as “cat” or at a superordinate level as “animal” or at various subordinate levels below the basic level as “Persian cat” or “Felis silvestris catus longhair Persian.”

In [11], the authors of Connotea provide a hands-on description of tagging systems. The study includes a snapshot of the tagging systems available, as of early 2005, and a breakdown of the key technologies behind these systems into a two-dimensional taxonomy. The first two facets of the first dimension in their taxonomy represent the identity of taggers: “tag user” and “content creator”. Both facets can be classified in a second dimension as either “self” or “others”. Other categorizations that the authors offer divide the space of tagging systems according to the “audience” (scholarly or general) and the “type of object store in the system” (URLs versus actual content). The same authors describe their own system—a social tagging system for academic articles—in a second article [12]; the technological and interaction techniques are described in depth, and an initial study of tag distribution is offered. The taxonomy we provide in Section 4 will expand upon the dimensions noted in their classifications.

Inherent in our model of tagging systems are connections or links between resources. As mentioned above, research on link-based systems in the context of the web is hardly new [1]. Obviously, the PageRank algorithm [18] had a significant impact on the field and on the way we use the web today, by supplying a mechanism to assess the importance of web pages. Lately, link analysis has been suggested to help fight web spam [10] by identifying trusted resources and propagating trust to resources that are linked from trusted resources. In tagging systems, similar concepts can utilize the information and trust in the social network and the links from users to resources (as well as between resources as before) to reason about the importance and trust of users *and* resources.

Perhaps more closely related to our tagging system model, Kleinberg [13] suggested an algorithm to identify web pages that are “hubs” and nodes that are “authorities” in a linked graph of resources, given a query term. In his model, Kleinberg views the hubs and authorities as a bi-partite graph, similar to the way we depict *users* and resources in our model in Figure 1. Taking the same hubs and authorities approach an inch closer to our model, Chakrabarti et al [5] extended Kleinberg’s work to include *anchor text*. Anchor text, the text that appears around a link to a certain resource, can be considered to have a similar role to tags in our model. Traditionally, the anchor text is associated with the resource the link is pointing to. The exact way the text is picked and associated with the resource varies between systems. Tags have the potential to increase comprehensiveness and accuracy of anchor-text based methods by treating the user and the resource separately in relevance metrics.

Also inherent in our model of tagging systems are relationships between users, a form typically described as a social network. While the social network literature related to tagging systems is too broad for the focus of this paper, we will summarize some of the important contributions. Social networks can be used both as a methodology for studying the social nature of tagging in these systems, as well as a tool for systems to expose relationships to users. A number of measures are applicable to each of these tasks, both from systemic and user-based perspectives. Centrality is a measure of how integral an individual is to a network [7], and can expose users whose social ties or tagging practices establish them as an influencer or opinion leader. Structural equivalence describes the similarity between two users based on the overlap in their personal networks [4], and can be used to find analogous users within the system. Partitioning a network into smaller structures can be helpful to both users and researchers; clustering addresses this problem by finding cohesive subgroups [25], while blockmodeling finds groups of users with similar roles within the network [26].

Like tagging systems, collaborative filtering (CF) is concerned with the relationships between people and resources, and the extent to which these connections can be leveraged to help users find new resources and people they would otherwise miss. Some of these systems have leveraged user-contributed metadata in the matching process, but this extra information is typically used as a filter after a match has been made [15]. To this extent, social tagging systems could be seen as complementary to CF, as tags are the primary means of finding similar resources; people have stipulated that these two systems would marry well, feeding each other with recommended content [21]. CF techniques have been studied extensively [3], and many are employed in popular tools, such as Amazon.com.

The research related to tagging systems separately covers each part of our model—people, resources, tags and the pairwise connections between them. To accurately describe the properties of systems including and connecting all of these components, we have integrated and extended background research for each of these components, spanning the fields of computer science, information science, and social networks. Each of these components is necessary to understand the relationship between objects, the words that describe them, and the motivations people have to do so. In the following section we will introduce a number of example tagging systems, followed by a descriptive taxonomy that shows how all of these pieces fit together in practice.

3. EXAMPLE TAGGING SYSTEMS

In this paper, we reference numerous tagging systems to show variations in architecture and incentives. We do not analyze most of these potentially ephemeral sites in depth, although we provide references to them in order to ground the reader with examples. For the sake of legibility, here is a brief description of sites we reference. There are many other tagging systems in existence, but we chose twelve that are representative of the diversity of those that are currently well used.

- *Del.icio.us* (<http://del.icio.us>): a “social bookmarking site,” allowing users to save and tag web pages and resources.
- *Yahoo! MyWeb2.0* (<http://myweb.yahoo.com>): similar to Del.icio.us, but including a social network of contacts.
- *CiteULike* (<http://www.citeulike.org/>): a site allowing users to tag citations and references, e.g. academic papers or books.
- *Flickr* (<http://www.flickr.com>): a photo sharing system allowing users to store and tag their personal photos, as well as maintain a network of contacts and tag others photos.
- *YouTube* (<http://www.youtube.com>): a video sharing system allowing users to upload video content and describe it with tags.
- *ESP Game* (<http://www.espgame.org/>) [23]: an internet game of tagging where users are randomly paired with each other, and try to guess tags the other would use when presented with a random photo.
- *Last.fm* (<http://www.last.fm>): a music information database allowing members to tag artists, albums, and songs
- *Yahoo! Podcasts* (<http://podcasts.yahoo.com/>): a site that indexes *podcasts* (regularly updated audio content), and allows users to tag them.
- *Odeo* (<http://www.odeo.com/>): another podcast information system supporting tagging and search.
- *Technorati* (<http://www.technorati.com/>): a weblog aggregator and search tool allowing blog authors to tag their posts.

- *LiveJournal* (<http://www.livejournal.com/>): a weblog and community website allowing users to tag their personal profile, along with individual blog posts
- *Upcoming* (<http://upcoming.org/>): a collaborative events database where users can enter future events (e.g., concerts, exhibits, plays, etc.) and tag them.

4. TAXONOMY OF TAGGING SYSTEMS

While we sometimes refer to social tagging systems as a coherent set of applications, it is clear that differences between tagging systems have a significant amount of influence on resultant tags and information dynamics. It is also clear that the personal and social incentives that prompt individuals to participate affect the system itself in various ways. We have developed two tagging taxonomies to analyze how 1) characteristics of system design and 2) user incentives and motivations may influence the resultant tags in a tagging system.

Different designs and user incentives can have a major influence on the usefulness of information for various purposes and applications, and in a reciprocal fashion, on how users appropriate and utilize these systems. The design of the system may solicit tagging useful for discovery, retrieval, remembrance, social interaction, or possibly, all of the above.

4.1 System Design and Attributes

We describe some key dimensions of tagging systems' design that may have immediate and considerable effect on the content and usefulness of tags generated by the system. For each dimension in our taxonomy, we note the ways in which the location of a system on this dimension may impact the behavior of the system. Some of these dimensions listed below interact; a decision along one of them may determine, or at least be correlated with, the system's placement in another.

- *Tagging Rights.* Possibly the most important characterization of a tagging system design is the system's restriction on group tagging. A tagging system can be restricted to *self-tagging*, where users only tag the resources they created (e.g., Technorati) or allow *free-for-all tagging*, where any user can tag any resource (e.g., Yahoo! Podcasts). This is not the apparent dichotomy that it seems, as systems can allow varying levels of compromise. For instance, systems can choose the resources users are to tag (such as images in the ESP Game) or specify different levels of permissions to tag (as with the friends, family, and contact distinctions in Flickr). Likewise, systems can determine who may remove a tag, whether no one (e.g., Yahoo! Podcasts), anyone (e.g., Odeo), the tag creator (e.g., Last.fm) or the resource owner (e.g., Flickr). The implication for the nature of the tags that emerge is that free-for-all systems are obviously broad, both in the magnitude of the group of tags assigned to a resource, and in the nature of the tags assigned. For instance, tags that are assigned to a photo may be radically divergent depending on whether the tagging is performed by the photographers, their friends, or strangers looking at their photos.
- *Tagging Support.* The mechanism of tag entry can have great impact on tagging system behavior. Observed systems fall into three distinct categories: *blind tagging*, where a tagging user cannot view tags assigned to the same resource by other users while tagging (e.g., Del.icio.us); *viewable tagging*, where the user can see the tags already associated with a resource (e.g., Yahoo! Podcasts); and *suggestive tagging*, where the system suggests possible tags to the user (e.g., Yahoo! MyWeb2.0). The suggested tags may be based on existing tags by the same user, tags assigned to the same resource by other users. Suggested tags can also be generated from or other sources of related tags such as automatically gathered contextual metadata, or machine-suggested tag synonyms. The implication of suggested tagging may be a quicker convergence to a folksonomy (see [9]). In other words, a suggestive system may help consolidate the tag usage for a resource, or in the system, much faster than a *blind tagging* system would. A convergent folksonomy is more likely to be generated when tagging is not blind. But it is not clear that consolidation is necessarily a good thing; arguably, a suggestive model may be applied carefully so that the agreement is not too widespread. As for viewable tagging, implications may be overweighting certain tags that were associated with the resource first, even if they would not have arisen otherwise.
- *Aggregation.* Another related feature of group dynamics comes from the aggregation of tags around a given resource. The system may allow for a multiplicity of tags for the same resource which may result in duplicate tags from different users; we term this approach the *bag-model* for tag entry (e.g., Del.icio.us). Alternatively, many systems ask the group to collectively tag an individual resource, thus denying any repetition; this interface we call a *set-model* approach for tag input (e.g., YouTube, Flickr). In the case that a bag-model is being used, the system is afforded the ability to use aggregate statistics for a given resource to present users with the collective opinions of the taggers; for instance, the tags around a popular link on Del.icio.us can be shown to the user to help characterize the breadth of opinions of the taggers. Furthermore, these data can be used to more accurately find relationships between users, tags, and resources given the added information of tag frequencies.
- *Type of object.* The type of resource being tagged is an important consideration. Sample objects types that are prominent in today's systems include, but are far from being restricted to, web pages (e.g., Del.icio.us, Yahoo! MyWeb2.0), bibliographic material (e.g., CiteULike), blog posts (e.g., Technorati, LiveJournal), images (e.g., Flickr, ESP Game), users (e.g., LiveJournal), video (YouTube) and audio objects such as songs (e.g., Last.fm) or podcasts (e.g., Yahoo! Podcasts, Odeo). In reality, any object that can be virtually represented can be tagged or used in a tagging system. For example, systems exist that let users tag physical locations or events (e.g., Upcoming). The implications for the nature of the resultant tags are numerous; a trivial example is that we suspect tags given to textual resources may differ from tags for resources/objects with no such textual representation, like images or audio, although this has not yet been empirically tested.
- *Source of material.* Resources to be tagged can be supplied by the participants (e.g., YouTube, Flickr, Technorati, Upcoming), by the system (e.g., ESP Game, Last.fm, Yahoo! Podcasts), or, alternatively, a system can be open for tagging of any web resource (e.g., Del.icio.us, Yahoo! MyWeb2.0). Some systems restrict the source through architecture (e.g., Flickr), while others restrict the source solely through social norms (e.g., CiteULike).
- *Resource connectivity.* Resources in the system can be linked to each other independent of the user tags. Connectivity can be roughly categorized as *linked*, *grouped*, or *none*. For

example, web pages are connected by directed links; Flickr photos can be assigned to groups; and events in Upcoming have connections based on the time, city and venue associated with the event. Implications for resultant tags and usefulness may include convergence on similar tags for connected resources, especially in *suggested* and *viewable* tagging support scenarios.

- *Social connectivity*. Some systems allow users within the system to be linked together. Like resource connectivity, the social connectivity could be defined as *linked*, *grouped*, or *none*. Many other dimensions are present in social networks, for example, whether links are typed (like in Flickr’s contacts/friends model) and whether links are directed, where a connection between users is not necessarily symmetric (in Flickr, for example, none of the link types is symmetric). Implications of social connectivity include, possibly, the adoption of localized folksonomies based on social structure in the system.

Table 1. Dimensions in the tagging system design taxonomy and possible implications

<i>Dimension</i>	<i>Main categories</i>	<i>Summary of Potential implications</i>
Tagging Rights	Self-tagging, permission-based, Free-for-all	Nature and type of resultant tags; role of tags in system
Tagging Support	Blind, suggested, viewable	Convergence on folksonomy or overweighting of tags
Aggregation model	Bag, set	Availability of aggregate statistics
Object type	Textual, non-textual	Nature and type of resultant tags
Source of material	User-contributed, system, global	Different incentives, nature and type of resultant tags
Resource connectivity	Links, groups, none	Convergence on similar tags for linked resources
Social connectivity	Links, groups, none	Convergence on localized folksonomy

The design options taxonomy for tagging systems is summarized in Table 1, including a brief summary of the potential impact of the design choices on the resultant tags and the type of benefits that can be derived from the system.

4.2 User Incentives

Incentives and motivations for users also play a significant role in affecting the tags that emerge from social tagging systems. Users are motivated both by personal needs and sociable interests. The motivations of some users stem from a prescribed purpose, while other users consciously repurpose available systems to meet their own needs or desires, and still others seek to contribute to a collective process. A large part of the motivations and influences of tagging system users is determined by the system design and the method by which they are exposed to inherent tagging practices. While tagging has the potential to be valuable for numerous applications, users can be unaware of or uninterested in the broader design motivations; they might instead be persuaded

by the norms of their friends and how they think that a particular system fits into their use.

Tagging can be a public and sociable activity, but not all tags emerge with an intended audience. Many users begin with the conception that they are tagging for themselves; some begin to appreciate the sociable aspects over time, while others have no interest in that component. Since user incentives are influenced by the design of a given system, the motivations underlying tagging vary both by people and by systems.

Evaluating these practices requires an understanding of why people contribute and the resulting effects on output and performance of the tagging system. In this section we will articulate the various incentives that can be outwardly observed in current social tagging systems and show how they can influence the use and utility of tags.

The motivations to tag can be categorized into two high-level practices: *organizational* and *social*. The first arises from the use of tagging as an alternative to structured filing; users motivated by this task may attempt to develop a personal standard and use common tags created by others. The latter expresses the communicative nature of tagging, wherein users attempt to express themselves, their opinions, and specific qualities of the resources through the tags they choose.

Both of these practices differ based on intended audience and future expectation of use. The following list of incentives expresses the range of potential motivations that influence tagging behavior. They are not intended to be mutually exclusive; instead we expect that most users are motivated by a number of them simultaneously.

- *Future retrieval*: to mark items for personal retrieval of either the individual resource or the resultant collection of clustered resources (examples: tagging a group of papers on Del.icio.us in preparation for writing a book, tagging songs on Last.FM to create an adhoc playlist, tagging Flickr photos ‘home’ to be able to find all photos taken at home later). These tags may also be used to incite an activity or act as reminders to oneself or others (e.g., the “to read” tag). These descriptive tags are exceptionally helpful in providing metadata about objects that have no other tags associated.
- *Contribution and sharing*: to add to conceptual clusters for the value of either known or unknown audiences. (Examples: tag vacation websites for a partner, contribute concert photos and identifying tags to Flickr for anyone who attended the show).
- *Attract Attention*: to get people to look at one’s own resource because they are common tags. When “tag clouds” or other such lists that reflect popularity of tags are visible in the system, users may be incentivized to contribute tags that might affect that global view (and even to create spam tags.).
- *Play and Competition*: to produce tags based on an internal or external set of rules. In some cases, the system devises the rules such as the ESP Game’s incentive to tag what others might also tag. In others, groups develop their own rules to engage in the system such as when groups seek out all items with a particular feature and tag their existence. Some users take advantage of what is available and try to alter the system in the way they see fit. Knowing that tags appeared in a tag cloud based on the frequency of a given tag for a podcast, Odeo users attempted to construct sentences by adding and removing tags to change the order of the tags in the interface.

- *Self Presentation*: to write a user's own identity into the system as a way of leaving their mark on a particular resource. (for example, the "seen live" tag in Last.FM marks an individual's identity or personal relation to the resource.)
- *Opinion Expression*: to convey value judgments that they wish to share with others (for example, the "elitist" tag in Yahoo!'s Podcast system is utilized by some users to convey an opinion.)

This range of motivations in turn affects the types of tags that are produced for a given resource. Golder and Huberman have outlined 7 individual types of tags observed in their study of Del.icio.us [9]. The first five types they mention roughly identify properties of the objects, such as the source, attributes, category membership or qualitative properties; these tags could arise from organizational motivations, social ones, or both depending on the perceived audience. The sixth tag type, self-reference (e.g., *mystuff* or *mywork*), reflects a probable intent to communicate this ownership to an outside audience, or alternatively to be used for personal organization. The final type, *task-organization* (e.g. *toread* or *jobsearch*) suggests an intent for personal organization.

The architecture of a social tagging system reflected by the taxonomy provided in Section 4.1 does not explicitly affect the *type* of tag that users produce; instead, the design may influence the incentives that drive individuals to use the system. The types of tags observed can be seen as a resulting artifact of the different forms of motivation expressed through the resulting interaction.

5. Case Study: Flickr

Due to their popularity, social tagging systems have grown to cover a wide range of resources and communities, spanning the entire range of incentives described in the previous section. Instead of simply classifying a long list of potentially ephemeral tools, we will give a complementary example to those provided in previous work. The system we have chosen to investigate is Flickr, a popular photo-sharing site that considers tags as a core element to the sharing, retrieval, navigation, and discovery of user-contributed images. Flickr allows users to upload their personal photos to be stored online, but unlike other online photo tools, Flickr makes these photos publicly viewable and easily discoverable by default. This design decision, along with the emphasis on tagging, has allowed the site to expand quite rapidly over its short lifespan.

This growth has in part been due to the wide array of social interactions Flickr supports: in addition to uploading photos, users can also create networks of friends, join groups, send messages to other users, comment on photos, tag photos, choose their favorite photos, and so on. This abundance of communication tools and forms of social organization creates a highly interconnected media ecology that can lead users to distant people and places with only a few clicks. Tags are an important part of this environment, where tags act as a primary navigational tool for finding similar resources and people.

As previously noted, the most extensive analysis of a tagging system has been completed on data collected from the social bookmarking site Del.icio.us [9]. We have chosen Flickr to provide an alternate interpretation to the conclusions derived from this study. In nearly every category within our system taxonomy, Flickr occupies an alternative space from Del.icio.us: it contains *user-contributed* resources as opposed to *global*; tagging rights are restricted to *self-tagging* (and at best *permission-based*, although in practice self-tagging is most prevalent) instead of a *free-for-all*; tags are aggregated in *sets* instead of *bags*; and

finally, the interface mostly affords for *blind-tagging* instead of *suggested-tagging*.

These design decisions shape the incentive structures that drive people to tag resources. Since Del.icio.us is largely task-focused, namely storing bookmarks for future retrieval, organizational motivations are most dominant. While the social element of tagging is evident from the leveraging of the community contribution, a lack of communication systems (e.g. messaging or explicit social networks) deemphasizes non-organizational social incentives.

Flickr users, on the other hand, are also likely to tag for their own retrieval, but coupled with an abundance of communication mechanisms, the system design encourages gaming and exploration of tag use. Users are primarily motivated by social incentives, including the opportunities to share and play.

In the following analysis we present a preliminary analysis of tag usage within Flickr. We have had the opportunity to work directly with a subset of the database used by Flickr, specifically information about photos, tags, and the explicit social relationships between users (i.e., the "contact" network). Because our focus is on the usage of tags, we have selected only those users who have utilized this feature (i.e., used at least one tag to describe a photo) and only those photos that have had at least one tag applied. Of the millions of Flickr users, we have randomly selected a set of 25,000 for our analysis of individual behaviors; for the more complicated case of network analysis, we have chosen a further subset of 2,500.

This study is only a preliminary look at the dynamics of the Flickr system and is meant to expose interesting trends and topics in the Flickr data. These topics illustrate various aspects of tagging systems and their incentive structure, but we do not attempt to prove or assert any general conclusions about all tagging systems.

5.1 Tag Usage

Tags are not mandatory in the Flickr usage model. Within a social tagging system, tags are typically an optional feature in a larger resource organization task. Like Del.icio.us, the Flickr interface prompts users for metadata about each resource identified: a title, a caption, and a list of tags. In the case of both systems, the tag input comes third in the input interface, but also differentiates them from other resource management tools.

In addition to tagging one's own photos, the Flickr system also allows users to tag their friends' photos. However, this feature is not largely used; of the 58 million tags we have observed, only a small subset are of this type; an overwhelming majority of tags are applied by the owners of photos.

Tag usage patterns vary quite drastically among Flickr users, and as expected, so does the adoption of tagging behavior. Figure 2 shows the cumulative distribution function (CDF) for tag vocabulary size across the set of users. The value at a given value is the probability (Y-axis) that a random user has a set of distinct tags (X-axis) that is larger than that collection size. For example, the probability that a Flickr user has more than 750 distinct tags is roughly 0.1%. This distribution illustrates the fact that most users have very few distinct tags while a small group has extremely large sets of tags.

The relationship between tag usage and other types of input can be a good indicator of how useful or important users believe tags are to the experience of using the system. Within Del.icio.us, Golder and Huberman found that there was not a strong association between the number of bookmarks made and the

number of tags used to annotate those bookmarks [9]. We studied three activities within the Flickr environment: the number of uploaded photos, the count of the user's distinct tags, and the number of contacts designated by the user. For example, a certain user can have 100 photos with a total of 200 distinct tags across these photos, and be connected to 50 different contacts.

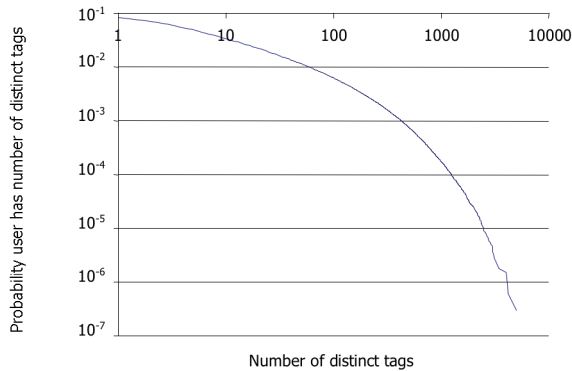


Figure 2. Distribution of distinct tag collections, represented as the probability that a r

Table 2 shows the pair-wise Pearson correlation [19] between photo collection size, distinct tags and number of contacts across the set of users. We computed this correlation for a set of 25,000 users randomly selected from our dataset. For example, the correlation between tags and photos is 0.518, suggesting a strong linear relationship between these variables, i.e. an increase in photo collection size implies an increase in the number of distinct tags. The strongest relationship between these three items (photos, distinct tags, and contacts) comes between photos and distinct tags, a likely relationship due to the fact that tagging ones' own photos is the dominant form of tags. The association between contacts and photos is much weaker than the one between contacts and distinct tags, which might suggest that tagging is related to social activity to some degree.

Table 2. Flickr usage correlation

	Tags	Photos	Contacts
Tags	1	.518	.386
Photos	.518	1	.192
Contacts	.386	.192	1

* $N = 25,000$

** $p < 0.001$ for all values.

In addition to social implications, another feature of tags worth investigating is an individual's use of tags over time. How does the frequency of tags change as a user becomes acclimated to the system? Do her tags become a cohesive taxonomy over time, or do they continue to grow as her experiences change? In studying Del.icio.us, Golder and Huberman show examples as to how certain users' sets of distinct tags continue to grow linearly as new resources are added. At the same time, they claim that the tags for a given resource tend to stabilize after only a few users have tagged it [9]. Since Flickr uses a set-model for representing tags, we cannot reexamine the latter observation, but we can look at the growth of a user's tags over time.

Figure 3 shows the growth of distinct tags for 10 randomly selected users over the course of uploaded photos. The users were selected as both frequent uploaders (greater than 100 photos) and

frequent taggers (greater than 100 tags). Each point on this graph shows the number of distinct tags (Y-axis) for a given user after the given photo number (X-axis). It is apparent from this plot that a number of different behaviors emerge from this social tagging system. In some cases (such as user A in Figure 3), new tags are added consistently as photos are uploaded, suggesting a supply of fresh vocabulary and constant incentive for using tags. Sometimes only a few tags are used initially with a sudden growth spurt later on, suggesting that the user either discovered tags or found new incentives for using them, as with user B. For many users, such as those with few distinct tags in the graph, distinct tag growth declines over time, indicating either agreement on the tag vocabulary, or diminishing returns on their usage. Despite the heavy usage of tags for each of the individuals whose tags are depicted in the figure, a number of classes of behavior have arisen, implying that the interaction between user, tag, and utility is a varied one.

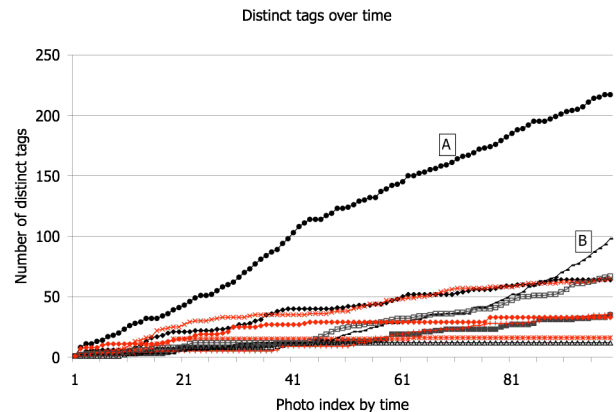


Figure 3. Number of distinct tags at given points in 10 random users' collections

Whereas Golder highlighted one form of tag vocabulary growth, namely growing at a diminishing rate over time, the graph illustrates two additional use classes each with several possible explanations. Is the case of linear growth related to the type of media being tagged, namely photos that are taken of constantly evolving subject matter? Or does it evolve from a motivation to continually attract new individuals to the users' photos? Likewise, the case of gradual increase could reflect a change in personal motivations (e.g., a need to start organizing photos once the collection grows above a certain size), or a social one (e.g., a sudden realization that tags can bring new people to see one's photos). These questions could be answered by looking at the relationship between the growth of users' tag collections and various forms of participation, such as the popularity of their photos or their use of the social network system.

5.2 Vocabulary Formation

All of the tagging systems we have mentioned in this paper are arguably social in nature; in some cases the social aspect comes from leveraging the community's collective intelligence, and in others there is explicit social interaction around the use of tags. Because Flickr allows users to enumerate social networks and develop communities of interest, there is a huge potential for social influence in the development of tag vocabularies.

One feature of the contact network is a user's ability to easily follow the photos being uploaded by their friends. This provides a continuous awareness of the photographic activity of their Flickr contacts, and by transitivity, a constant exposure to tagging

practices. Do these relationships affect the formation of tag vocabularies, or are individuals guided by other stimuli? To expand on this question, we have randomly chosen 2500 users with a considerable number of tags (greater than 100) and paired them with two other individuals: one randomly chosen from the rest of the set, and the other from their list of contacts. From these pairings we have calculated the overlap in their tag sets; the overlap is computed as $|A \cap B| / |A \cup B|$, where A and B are the sets of tags from our two users.

The results of this inquiry are depicted in Figure 4. This graph shows two frequency distributions for the overlaps between sets of users: the overlap between the given user and another randomly chosen one, shown with a dashed blue line, and the overlap between the same user and one of their contacts, shown with a solid red line. The random users are much more likely to have a smaller overlap in common tags, while contacts are more distributed, and have a higher overall mean.

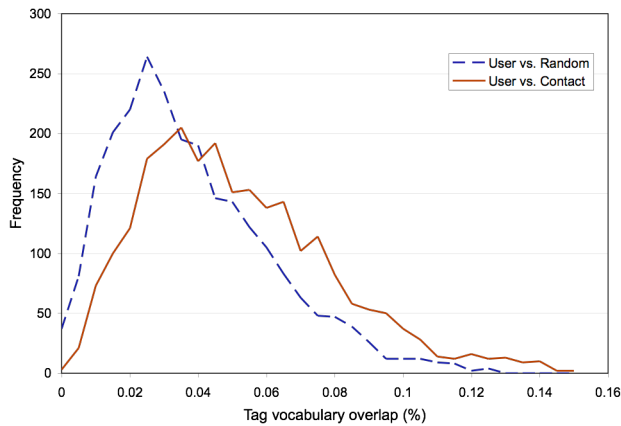


Figure 4. Vocabulary overlap distribution for random users and contacts (n=2500)

This result, while still preliminary, shows a relationship between social affiliation and tag vocabulary formation and use even though the photos may be of completely different subject matter. This commonality could arise from similar descriptive tags (e.g., bright, contrast, black and white, or other photo features), similar content (photos taken on the same vacation), or similar subjects (co-occurring friends and family), each suggesting different modes of diffusion.

Other likely explanations for the observed correlation between social connection and common tag usage may be found in the descriptive categories of sociolinguistics which studies how different geographic and social formations structure the coherence and diffusion of semantic and syntactic structures in various "lects" within a larger sociolinguistic system. Some of these example lects include: dialect (a lect used by a *geographically* defined community); sociolect (a lect used by a *socially* defined community); ethnolect (a lect spoken by a particular ethnic group); ecolect (a lect spoken within a household or family); and idiolect (a lect particular to a certain person). If we conceptualize social tagging systems within the theoretical frame of sociolinguistics, these and other "lects" seem especially applicable to understanding and classifying the apparent isomorphism between social and linguistic structures we observed in Flickr. The structures, changes, and diffusion within and amongst various "lects" in social tagging systems will likely have similar patterns to those found in social network analyses and in sociolinguistic language maps. Considering these sociolinguistic categories as

we attempt to compute structural isomorphism and the interactions between social structures and tagging structures (for example, hubs, bridges, and diffusion) may prove exceptionally useful in explaining the formation, efficacy, and dynamics of social tagging systems.

These questions call for a much deeper investigation of this phenomenon, a study that could answer many questions about the relationship between people, communication, and the emergence of common lects in social tagging systems.

6. CONCLUSIONS

Social tagging systems have the potential to improve on traditional solutions to many well-studied web and information systems problems. Such problems include personalized or biased link analysis, organizing information, identifying synonyms and homonyms, building networks of trust to combat link spam, monitoring trends and drift in information systems and more. The prospects of reasoning about tags, users, and resources in unity are encouraging.

In order to study these systems, researchers should observe the system's place within the taxonomy of architectures described in Section 3.1. Studies should also consider the incentives driving participation, and the extent to which the system supports or restrains these motivations. In studying Flickr, we showed that the dynamics of interaction and participation are different than those of Del.icio.us. Indeed, Flickr and Del.icio.us are rather distinct when positioning them in the dimensions of our taxonomy. Del.icio.us is a free-for-all, suggestive, bag-model (to mention just three key dimensions) system. Del.icio.us is therefore likely to generate a different use model and output than Flickr, a (mostly) self-tagging, viewable, set-model system. Moreover, the incentive models of Flickr and Del.icio.us are also substantially disparate, suggesting even more expected differences in the systems' output.

We hope that system designers will consider these design decisions in architecting their tagging systems. By laying out the implications of the choices in each dimension of our hierarchy, we hope to assist planners as well as researchers and academics.

Finally, by no means do we contend that the design taxonomy and incentive taxonomy we describe are complete. New uses for tagging systems are invented every day; users of such systems appropriate them with an ever-changing set of goals, motives, and aspirations. We hope that our taxonomy can serve as a foundation for researchers and enable a more complete understanding of the constraints and affordances of tag-based information systems.

7. ACKNOWLEDGMENTS

The authors would like to thank the members of the Flickr team, and the users of Flickr for providing us with fascinating data to study.

8. REFERENCES

- [1] Baeza-Yates, R. and Ribeiro-Neto, B.. Modern Information Retrieval. Addison-Wesley, 1999.
- [2] Brieger, R.L., 1991. Explorations in Structural Analysis: Dual and Multiple Networks of Social Structure. New York: Garland Press.
- [3] Breese, J.S., Heckerman, D. and Kadie, C.M. Empirical analysis of predictive algorithms for collaborative filtering. Microsoft Research Technical Report, (MSR-TR-98-12), October 1998.

- [4] Burt, R. 1992. *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.
- [5] Chakrabarti, S., Dom, B., Raghavan, P., Rajagopalan, S., Gibson, D., and Kleinberg, J. 1998. Automatic resource compilation by analyzing hyperlink structure and associated text. In *Proceedings of the Seventh international Conference on World Wide Web 7* (Brisbane, Australia).
- [6] Coates, T. Two cultures of fauxonomies collide. June 4 2005. http://www.plasticbag.org/archives/2005/06/two_cultures_of_fauxonomies_collide.shtml
- [7] Freeman, L. C. 1979. Centrality in Social Networks: Conceptual Clarification. *Social Networks*. 1, 215-239
- [8] Furnas, G. W., Landauer, T. K., Gomez, L. M., and Dumais, S. T. The vocabulary problem in human-system communication. *Commun. ACM* 30, 11 (1987).
- [9] Golder, S., and Huberman, B. A. The Structure of Collaborative Tagging Systems. HP Labs technical report, 2005. Available from <http://www.hpl.hp.com/research/idl/papers/tags/>
- [10] Gyongi, Z., Garcia-Molina, H., Pederson, J. Combating spam with trustrank. In *Proceedings of the 30th International Conference on Very Large Databases (VLDB)*, 2004.
- [11] Hammond, T., Hannay, T., Lund, B. and Scott, J. Social Bookmarking Tools – A General Overview. *D-Lib Magazine* 11, 4 (April 2005)
- [12] Hammond, T., Hannay, T., Lund, B. and Scott, J. Social Bookmarking Tools – A Case Study. *D-Lib Magazine* 11, 4 (April 2005)
- [13] Kleinberg, J. M. 1998. Authoritative sources in a hyperlinked environment. In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms* (San Francisco, 1998).
- [14] Lakoff, G. *Women, Fire and Dangerous Things*. University of Chicago Press, Chicago, 2005.
- [15] Malz, D. and Ehrlich, K. Pointing the way: Active collaborative filtering. In *the Proceedings of CHI 1995*.
- [16] Mathes, A. *Folksonomies – Cooperative Classification and Communication Through Shared Metadata*. UIC Technical Report, 2004.
- [17] Merholz, P. Clay Shirky's Viewpoints are Overrated. <http://www.peterme.com/archives/000558.html>
- [18] Page, L., Brin, S., Motwani, R. and Winograd, T.. *The PageRank citation ranking: Bringing order to the web*. Technical report, Stanford University, 1998.
- [19] Rice, J.A., *Mathematical statistics and data analysis*. Belmont, CA: Duxbury Press (1995)
- [20] Shirky, C. *Ontology is Overrated: Categories, Links, and Tags*. http://shirky.com/writings/ontology_overrated.html
- [21] Udell, Jon. Collaborative filtering with Del.icio.us. June 23, 2005. <http://weblog.inworld.com/udell/2005/06/23.html>
- [22] Vander Wal, T. *Folksonomy Definition and Wikipedia*. November 2, 2005. <http://www.vanderwal.net/random/entrysel.php?blog=1750>
- [23] von Ahn, L. and Dabbish, L. 2004. Labeling images with a computer game. *CHI 2004* (Vienna, Apr. 2004). ACM Press, 319-326.
- [24] Walker, J. Feral hypertext: when hypertext literature escapes control. In *Proceedings of the Sixteenth ACM Conference on Hypertext and Hypermedia* (Salzburg, Austria, Sept. 2005). *HYPERTEXT '05*. ACM Press, New York, NY, 46-53.
- [25] Wasserman, S. and Faust, K.. *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press, 1994.
- [26] White, H.C., Boorman, S.A., and Breiger, R.L. 1976. Social structure from multiple networks: Blockmodels of roles and positions. *American Journal of Sociology*. 81, 730-779