

General patterns of tag usage among university groups in Flickr

Item Type	Journal article
Authors	Angus, E;Thelwall, Mike;Stuart, David
Citation	Angus, E., Thelwall, M. and Stuart, D. (2008) General patterns of tag usage among university groups in Flickr, Online Information Review, 32(1), pp. 89-101, https://doi.org/10.1108/14684520810866001
DOI	10.1108/14684520810866001
Publisher	Emerald
Journal	Online Information Review
Download date	2025-05-24 05:47:39
License	https://creativecommons.org/licenses/by-nc-nd/4.0/
Link to Item	http://hdl.handle.net/2436/622430

General patterns of tag usage among university groups in Flickr

Emma Angus, Mike Thelwall & David Stuart

Purpose - There has been a significant rise in the use of web 2.0 social network websites and online applications in recent years. One of the most popular is Flickr: an online image management application. This paper investigates general patterns of tag usage and determines the usefulness of the tags used within university image groups to the wider Flickr community.

Design/Methodology/Approach - This study uses a webometric data collection, classification and informetric analysis.

Findings - The results show that members of university image groups tend to tag in a manner that is of use to users of the system as a whole rather than merely for the tag creator.

Originality/Value - This paper gives a valuable insight into the tagging practices of image groups in Flickr.

Keywords - Webometrics; classification; universities; Flickr; images; tagging.

Paper Type – Research paper.

Introduction

In recent years there has been a significant rise in social network websites and online applications where, 'like minded users share resources, create, tag and label content and rate it in some way' (Matthews, 2006), so-called 'web 2.0' applications (O'Reilly, 2005). Such applications vary greatly, ranging from websites where you can add, organise and share: bookmarks (e.g., del.icio.us), academic references (e.g., CiteULike.org), and photographs (e.g., Flickr.com). One thing that all of these websites have in common is their emphasis on online collaboration and the sharing of resources among users.

Although social websites have become very popular with users and have received a lot of media attention (e.g., BBC News, 2006) there has been relatively little research into the users of such sites and what they are being used for. It seems as though they have an important role to play in enabling people to utilise and share pertinent information, however this is an area that needs further research. This study investigates one of the popular new social websites, Flickr, an online image management and sharing application; it explores what image tags reveal about the purpose of Flickr for the image uploader, and whether the tags are useful for the subsequent searching for and retrieval of images.

Flickr allows users to upload, store and share digital images which may be 'tagged' (i.e., assigned one or more freely chosen keyword(s)) to aid in their organisation and retrieval. Such images can be made private, viewable to selected friends and family only, or they can be viewable to the entire user community of the system. Users of Flickr may also choose to be part of a group(s). In groups, 'like-minded users gather, discuss things, and share pictures' (Wilkinson, 2007). The theme or topic of a group can be based around anything at all, ranging from a particular place/holiday destination, to the 'squared circle' group, where users upload photographs which must be of equal height and width, and contain a circle. Groups can either be public, public (invitation only) or completely private (Bausch & Bumgardner, 2006). Public (invitation only) groups are those where only people who are invited can gain access or upload images to that particular group. For users of Flickr who make their images publicly available to the entire user community of the system, and particularly those images that are placed in a public 'group', it seems reasonable to assume that they want their images to be found and viewed and are not merely using the application as a place to store and organise images purely for their own benefit.

Related Literature

Tagging

Literature published to date in the field of web 2.0 tagging has focussed on a number of key areas such as: providing an overview of the emergence of tagging and reviewing related literature (Hammond et al., 2005; Marlow et al., 2006; Macgregor & McCulloch, 2006; Shirky, 2005); the advantages and

disadvantages of tagging and how it resembles/differs from traditional classification (Kipp & Campbell, 2006; Macgregor & McCulloch, 2006; Mathes, 2004; Shirky, 2005); recommendations for improving the quality of tags (Guy & Tonkin, 2006); tag folksonomy (Guy & Tonkin, 2006; Mathes, 2004; Quintarelli, 2005); comparing structured and unstructured tagging (Bar-Ilan et al., 2006); and looking at tag usage patterns (Golder & Huberman, 2006; Guy & Tonkin, 2006; Kipp & Campbell, 2006; Marlow et al., 2006) However, studies of tag classification and usage patterns have tended to investigate the tags used for text-based resources rather than for images, and so have limited value in this investigation.

The term folksonomy, a portmanteau of the words 'folk' and 'taxonomy', was coined by Thomas Vander Wal (Smith, 2004) to describe the emergent practice of tagging whereby, 'people use their own vocabulary in order to add explicit meaning to the information or object they are consuming' (Anderson, 2007). Vander Wal describes this as a 'bottom-up social classification'. Quintarelli (2005) adds to this by describing it as a 'grassroots' approach to classification where people are moving away from hierarchical authoritative schemes. Vander Wal (2005) defines two typologies of folksonomy: broad and narrow. Broad folksonomies are defined as those where many people tag the same item (as in the Del.icio.us system where multiple users tag the same URLs). Narrow folksonomies (as in the Flickr system) are defined as those where users are more likely to be tagging content for their own retrieval purposes (i.e., their own collection of uploaded images). However, in the vast majority of web 2.0 sites, tags assigned by individual users to resources are visible to the rest of the user community of the system regardless of whether or not they fall under the definition of a broad or narrow folksonomy; meaning that they both have the potential to persuade other users to use the same tags and foster consensus. Mathes (2004) hypothesised that tag usage in web 2.0 systems would follow a power law scenario, whereby the most used tags are more likely to be used by other users since they are more likely to be seen. Guy & Tonkin (2006) tested this hypothesis on tag usage in the Flickr (narrow folksonomy) and Del.icio.us (broad folksonomy) systems and found that tag usage followed a power law distribution in both instances.

In an analysis of the usage patterns of the del.icio.us collaborative bookmarking and tagging system, Golder and Huberman (2006) suggest that a significant amount of tagging, if not all, is done for personal benefit. Such conclusions, however, are based on the frequency distribution of tag usage; they deem that the tags that are used most frequently by users are the tags which are generally most 'useful'. Although they did not include a classification of the tags they looked at, they stated what they believed were the functions of tags to perform:

1. Identifying what (or who) the bookmarked page is about.
2. Identifying what it is.
3. Identifying who owns it.
4. Refining categories (refine or supplement existing categories).
5. Identifying qualities or characteristics (adjectives such as 'funny' or 'inspirational').
6. Self-reference (MyStuff, MyWebsite).
7. Task-organising (ToRead, ToPrint).

Golder and Huberman (2006) define categories 1 - 4 as tags that are 'extrinsic' to the tagger, and categories 5 - 7 as those that are 'relative to or only relevant to the tagger'.

Golder and Huberman's (2006) notions of tags being extrinsic to the tagger or relative only to the tagger may be considered broadly similar to the Marlow et al. (2006) premise that motivations to tag can be categorised into two main practices: organisational and social. Organisational practices arise from the use of tagging as an alternative to structured electronic filing (i.e., filing items by date, type etc. for personal future retrieval). Social practices express the communicative nature of tagging, where people tag so that other users of the system are able to view and retrieve their content. Similarly, Hammond et al. (2005) define these user motivations as 'selfish' and 'altruistic' tagging disciplines. Hammond et al. argue that the nature of a web application is responsible for driving a particular tagging practice for its users. They claim that because Flickr users are likely to be managing

personal collections of their own photos, they are far more likely to adopt a ‘selfish’ tagging discipline. In contrast, Marlow et al. argue that users in Flickr are primarily motivated by social incentives to tag, including the opportunity to share and view other users’ images.

Image Meaning and Indexing

In a classification of the tags used for visual resources, it is useful to consider the work of image interpretation and theory, and image indexing, rather than focussing solely on the work of tag usage in relation to web 2.0 systems.

The art historian Erwin Panofsky developed the theory of ‘three levels’ of meaning in a work of art (Panofsky, 1983) and his work forms the basis for much of the subsequent theoretical work on the classification of images (Jørgensen, 2003). When looking at images, the primary (pre-iconographic) level of interpretation and description will be based on the recognition of objects within the image and factual information, no subject specific knowledge is needed to make this description. The secondary (iconographic) level involves an interpretation of the objects in the picture and requires a familiarity with the cultural context of the image. The third (iconological) level of description attempts to determine the intrinsic meaning of the content of the image.

Panofsky’s theory of image meaning is reiterated in Shatford’s (1986, 1994) work on image indexing. Shatford argues that regardless of what type of image you are looking at, all images have attributes which can be categorised and generalised. Specifically, there are four general categories which attributes can fall into:

- Biographical attributes (i.e., who created the image; who it belongs to; the time and date it was created).
- Subject attributes (i.e., images may be both *of* and *about* something - what an image is *of* is likely to be concrete and objective – pre-iconographic; what an image is *about* is likely to be abstract and subjective - iconographic; images can also be both generic and specific at the same time - whilst an image may be specifically useful for its depiction of the Clifton suspension bridge in Bristol, UK, it may also be generically useful for its depiction of a bridge).
- Exemplified attributes (i.e., an image may *be* an etching or a photograph, which is different to an image that is *of* an etching or a photograph).
- Relationship attributes (i.e., images can be related to or associated with other images or textual works, for example illustrations which accompany text).

Research Questions

Whether Flickr image tags tend to be for personal organisation or for the benefit of wider users is likely to vary and will be dependent upon the context in which the images are placed. With such a broad range of users utilising Flickr, this pilot study focuses upon one particular category of image group: universities. This category was chosen as the subject is familiar to the researchers. Although it is not clear how much of the participation in university groups is on a formal level (e.g., university marketing departments) or on an informal level (e.g., individual students), these groups have the potential to be of great value to universities, providing an instant visual entry point into the world of university life and could therefore be used as a valuable marketing tool for attracting potential students and collaborating institutions. However, such practices would require the images to be tagged in a useful manner. Using a webometric data collection, classification and informetric analysis this study investigates the tagging practices in university groups and addresses the following question: *Are the images in university groups tagged in a manner that is useful for people other than the person who uploaded the image?*

More specifically, it aims to answer the following subset of questions relating to tag usage patterns:

- What function do tags generally perform in Flickr?
- Do image tags tend to follow a power law distribution?
- How exhaustive are Flickr users in assigning tags to the images they upload?

Methods

Based on the research questions identified, it was decided that the investigation would be carried out in three main parts: data collection; a classification of image tags (in order to determine the general function of the sample of tags); and finally, an analysis of the general patterns of tag usage using informetric techniques (specifically looking at the number of tags used per image and comparing tag frequency to a normal Zipfian distribution).

Data Collection

Utilising Flickr's Application Programming Interface (API) a programme was written which retrieved a list of all of the groups within Flickr which had the term 'university' as part of their group name or description (e.g., Wolverhampton University). Ten of the most recently uploaded images from each of the university groups was retrieved, giving a set that was not biased in favour of the groups with larger image collections. Where this was not possible, due to there not being ten publicly available images, as many as possible were retrieved. The selection of images generated were then assigned with a random number, and from this, a sample of 250 random images were taken, along with all of their accompanying tags.

Classification

All of the 250 images were looked at alongside their tags, and tags were classified according to their perceived relationship with their accompanying image. It was decided that tags could not be looked at in isolation due to the complex nature of image interpretation. Golder and Huberman (2006) point out the problem with polysemous words and tags. Polysemous words have many related senses, for example, 'a 'window' may refer to a hole in the wall or to the pane of glass that resides within it' (Pustejovsky in: Golder & Huberman, 2006). Homonyms would also cause a problem if tags were looked at in isolation (i.e., the tag 'bat' could mean the small nocturnal mammal or a wooden implement with a handle). Without looking at the tag's accompanying image it would be impossible to distinguish the true meaning of the tag.

Drawing on the work of Panofsky (1983), Shatford (1986, 1994), and Golder and Huberman (2006), and also by visiting a random selection of tags and images, a classification scheme was developed which comprised of possible categories for all image tags in Flickr (see Table 1, column 1 in the results section). Using the three levels of meaning in a work of art as defined by Panofsky (1983), and the notions of 'of' and 'about' as defined by Shatford (1986, 1994), categories A1, B1a, B1b and B2 were developed. Drawing on the work of Golder and Huberman (2006), categories C1, C2, C3, and C4 were added. Categories C5, D1, D2 and D3 were created as a result of visiting a random selection of images and tags within the Flickr site. See Appendix 1 for the classification protocol and a detailed breakdown of the categories.

Mirroring the work of Golder and Huberman (2006), the classification scheme used in this study can arguably be split into tags which can be seen to be useful for the user community of the system as a whole, and tags which can be seen to be beneficial to only the tagger themselves (see Table 1, column 2 of the results section). This split, however, differs quite considerably from Golder and Huberman's distinctions. Whilst they define the tag category of 'identifying who owns' the resource as being extrinsic to the tagger, we believe that unless users of Flickr are specifically searching for images uploaded by people that they know, then a tag which denotes ownership is likely to be of no real value. Conversely, Golder and Huberman define the tag category of 'identifying qualities or characteristics' as being relevant only to the tagger. This tag category is essentially an iconographic level of interpretation, and is therefore useful to other users who are searching for images based on the same level of interpretation.

Category D3 (foreign words/characters) was not included in this split. The analysis was carried out from the perspective of an English speaker; therefore distinctions could not always be made regarding which category a foreign word/character should be assigned to.

Informetric Analysis

An informetric analysis was carried out into the frequency of tags used and the number of tags used per image for all of the images retrieved (i.e., up to ten per university group). The frequency of tags used gives an indication of the range of language (tags) used and whether this corpus is following a traditional language distribution (Boyce, Meadow & Kraft, 1994). It also allows further investigation into Mathes's (2004) hypothesis that tag usage will follow a power law distribution. Analysis of the number of tags used per image gives an indication of the effort put into the tagging process by image up-loaders.

Results

A total of 2,374 university groups were identified using the Flickr API, of which 1,983 had publicly accessible images. The initial sample of up to ten images from each group resulted in 15,353 images, of which 12,812 had tags assigned.

The classification sample of 250 random images had a total of 2,006 accompanying tags. Table 1 shows the total amount of tags that fell into each of the classification categories, as well as the percentage of overall tags that could be deemed as useful to the Flickr community as a whole, and those that would only be useful to the individual/group.

Table 1. Frequency and percentage of tag categories for 2,006 tags from random Flickr university image groups

Classification Category		Tag Motivation	Overall Percentage of Tags	Frequency & Percentage	
A1	Tag generically identifies what image is 'of'	Useful to Flickr community as a whole Social (Marlow et al., 2006) Altruistic (Hammond et al., 2005)	52%	309	15%
B1a	Tag specifically identifies what image is 'of' (place names/events)			289	14%
B1b	Tag specifically identifies what image is 'of' (people/animals)			43	2%
B2	Tag identifies what image is 'about'			398	21%
C1	Refining tag	Useful only to individual/group Organisational (Marlow et al., 2006) Selfish (Hammond et al., 2005)	39%	262	13%
C2	Self-reference tag			244	12%
C3	Task-organising tag			0	0%
C4	Tag which denotes ownership			31	2%
C5	Compound tag			233	12%
D1	Misspelling			5	0%
D2	Unable to determine relationship			5	0%
D3	Foreign word/character		9%	187	9%

The results show that over half of all the tags which were classified (excluding foreign words/characters) fell into categories which seem to be useful for the user community of the system as a whole.

The results from the classification exercise show that the highest proportion of tags fell into the B2 category – tags which identify what the image is about (e.g., adjectives and descriptive terms).

As Table 1, column 4 shows, 12% of all tags classified were compound tags (i.e., words/phrases/sentences joined together as one long text string). Whilst in the vast majority of cases it was obvious to the researcher when classifying these tags what the separated words would be, compound tags are of limited use as keyword entry points for subsequent users carrying out image searches. It should also be noted that Flickr does allow for spacing between terms in tag names, therefore tags without any spacing between terms seems unnecessary. Encouraging their replacement would, 'improve the potential for richer search functionality' (Guy & Tonkin, 2006).

Number of Tags Used

Altogether 14% of the images had no tags. For those images that did have tags assigned to them the mode average for the number of tags used was four, with 9% of all of the images looked at in the analysis having four tags assigned to them. A total of 53% of the images had more than four tags, which suggests that image uploaders put considerable effort into tagging their images; this should hopefully aid with subsequent image retrieval. It should also be noted that images which had a high number of tags assigned came from a variety of different university groups, indicating that no one particular university was overly dominant or exhaustive with their tagging. Overall it was found that two particular genre of images tended to dominate the university groups and these were: photographs of an architectural nature (e.g., university buildings) and photographs of groups of people (e.g., graduation images or groups of friends posing for the camera), and these two genres tended to attract an equal amount of tagging from users.

Frequency of Terms Used

A further analysis was carried out which looked at the tag frequency for the 15,353 images (i.e., up to ten for each of the groups with publicly accessible images). There were 114,678 tags used in total, with 26,599 unique tags.

With a logarithmic transformation applied to the axes, Figure 1 shows that the terms used within the Flickr university groups loosely follow a Zipfian power law distribution (Wolfram, 2003) whereby the frequency of any given term is inversely proportional to its popularity rank. The results indicate that single use tags do not dominate the Flickr system; instead, natural agreement regarding term usage tends to emerge through a bottom-up consensus (Wright, 2004) where many users agree on using a few popular tags. This is in support of Mathes's (2004) and Guy and Tonkin's (2006) research into tag popularity and distribution. However, the slightly skewed nature of the distribution can perhaps be attributed to the fact that the tags looked at were from a biased sample (i.e., university image groups) with the top tag, 'university' accounting for 0.95% of all tags used, along with '2006' (0.8%), with 'Architecture' (0.4%), 'College' (0.37%), and 'Campus' (0.35%) being the highest-ranked tags in the 'normal' part of the distribution.




Figure 1. Frequency of terms used in Flickr university image groups

Discussion

The results support Marlow et al.'s (2006) presumption that users of Flickr are primarily motivated by social factors when tagging images, rather than Golder and Huberman's (2006) and Hammond et al.'s (2005) claims that tagging is primarily for personal benefit. This is an important point to note as it provides further evidence in removing some of the speculation surrounding motivations to tag and the use of applications such as Flickr. The results from this investigation would suggest that Flickr users are primarily motivated to tag by social incentives, including the opportunity to share and view other

users' images (Marlow et al., 2006). Nevertheless, this result is partly attributable to the choices made when determining the usefulness of certain types of tags. However, it must also be noted that whilst organisational and selfish tags are only of real value to individuals/groups, social and altruistic tags can be of value to both to individuals/groups and to the wider Flickr community. So, whilst it can be presumed that users of Flickr are primarily motivated by social factors when tagging images, it could in fact be that they are tagging only for personal benefit whilst at the same time using tags which are social/altruistic in nature.

The classification scheme that was used for this study needs discussion because no research has been published to date which specifically looks at tag classification for images. Previous research has tended to focus on image interpretation and the meaning in visual content, and also the classification of tags for text-based resources. Whilst the work carried out by Shatford (1986, 1994) and Golder and Huberman (2006) gave a useful starting point for the development of the classification scheme, it has been found that the bespoke classification scheme used in this study was not a perfect match for tags used within Flickr. It is important in an investigation of images to take into consideration the concepts of image interpretation.

The fact that no tags were found to fall under the task-organising category (C3) can be attributed to the nature of content stored within Flickr (i.e., images rather than text-based resources). Perhaps there is not the same need for task organisation with images as there is with text as any 'actions' which they may have necessitated (i.e., altering of size, cropping, manipulation) will have presumably been carried out using image management software prior to being uploaded into Flickr for viewing and storing. Nevertheless this also suggests that Flickr publishing may be an end goal rather than an intermediate step (as is the case with some bookmarking websites).

Whilst every effort was made to ensure that the tag categories developed for the classification remained exclusive and independent of one another, many tags could lie between categories or equally well in multiple categories (e.g., an image tagged with the word 'Sarah' could be a photograph 'of' a girl named Sarah or 'belong' to a girl named Sarah). Looking at tags alongside their accompanying images was an attempt to minimise this variable. Nevertheless, the lines ultimately drawn are subjective.

Although Golder and Huberman (2006) assert that the most frequently occurring tags are likely to be those that are most generally 'useful', this is not necessarily true for Flickr. The frequency distribution of the tags shown in this study is too similar to the frequency distributions seen in natural language to necessarily discount the usefulness of infrequent tags. This point is further corroborated by research carried out by Goodrum and Spink (2001) into image searching on the web. They analysed a data set of 1,025,908 user queries from the Excite image search engine and found that the distribution of highly occurring search terms falls off very quickly when plotted in a log rank-frequency graph, with a very large number of search terms having a frequency of only one. All of the search terms analysed were obviously deemed to be 'useful' to the user who originally submitted them, regardless of whether or not the terms produced an appropriate selection of retrieved images. A further factor to take into account when looking at frequently occurring tags is the fact that not all groups included in this study may have actually been university groups. The fact that anyone is able to create a group on Flickr and call it whatever they like, means that tags from groups which have nothing to do with higher education institutes may have been analysed as part of this investigation. Whilst the top five tags would suggest that the majority of groups looked at were groups relating to universities it is an important variable to bear in mind nonetheless.

The number of tags used per image in university groups varied considerably, with 14% of images having no tags, whilst one user had assigned 629 tags to one particular image. This study has not made a comparison between the number of tags used per image and the usefulness of the accompanying tag set for that particular image, but it seems likely that there is a cut off point, beyond which the tags are no longer useful but rather reflect a quirk of the image uploader. That over half of the images looked at had more than four tags indicates that users are being fairly thorough in the tagging process. However,

no attempt was made in this study to investigate duplicate tagging. A Flickr user who uploads a large amount of images in a particular session could perhaps assign the same set of tags to all of the images uploaded in that particular session resulting in a large number of images all with a high number of tags. However, the tag sets assigned may not necessarily be appropriate for all of the images which were uploaded in that session meaning that the user was casual rather than thorough with their tagging, despite the large number of tags.

An important point to bear in mind is that tags looked at in this investigation came from images uploaded to university orientated image groups, it is therefore likely that the findings do not apply to other types of image groups within Flickr. For example, image groups oriented towards professional photographers may have a much higher average proportion of generally useful tags (A1-B2), whereas more personal groups may tend to have a higher proportion of more narrowly useful tags (C1-D1).

Conclusion

In answer to our research question, from the results presented in this investigation it seems that tags in university image groups tend to be useful for image retrieval purposes. In support of this claim, the majority of tags are of types that are useful to the Flickr community and image uploaders take care to assign multiple tags to their images, at least four for the majority of images. In addition, there is a broad and natural spread in the range and frequency of use of the individual tags, with the exception of the two most popular tags, 'university' and '2006'.

This study is merely a first step into investigating the potential of Flickr and social tagging generally. Future avenues of research include:

- A refinement of the classification scheme used for image tags, which draws more directly from the theory of image retrieval (e.g., Jörgensen, 2003).
- A refinement of the classification scheme used for image tags, which draws more directly from the field of image interpretation and theory (e.g., Panofsky, 1983; Shatford, 1986, 1994). This may be most relevant to groups with a specific artistic focus.
- Research into whether there are differences in the tagging practices of groups and individuals in order to test the hypothesis that people tagging as part of a group will use tags which are more social/altruistic.
- A comparison of the tagging practices of different types of groups to assess whether there is a general rule as to which types of groups tend to have the most generally useful tags.

References

- Anderson, P. (2007), "What is Web 2.0? Ideas, technologies and implications for education", *JISC Technology and Standards Watch*, February 2007, available at: <http://www.jisc.ac.uk/media/documents/techwatch/tsw0701b.pdf>
- Bar-Ilan, J., Shoham, S., Idan, A., Miller, Y., & Shachak, A. (2006), "Structured vs. unstructured tagging – A case study", *Proceedings of the 15th International World Wide Web Conference*, (Collaborative Web Tagging Workshop), May 22, Edinburgh, UK.
- Bausch, P. & Bumgardner, J. (2006), *Flickr Hacks: Tips and Tools for Sharing Photos Online*, Sebastopol, CA: O'Reilly.
- BBC News (2006), "Rise of the web's social network", *BBC Website*, available at: http://news.bbc.co.uk/1/hi/programmes/click_online/5391258.stm (accessed 27 November 2006).
- Boyce, B. R., Meadow, C. T. & Kraft, D. H. (1994), *Measurement in Information Science*, San Diego, CA: Academic Press.
- Golder, S. A. & Huberman, A. (2006), "Usage Patterns of Collaborative Tagging Systems", *Journal of Information Science*, Vol. 32 No. 2, pp. 198-208.
- Goodrum, A. & Spink, A. (2001), "Image searching on the Excite Web search engine", *Information Processing and Management*, Vol. 37 No. 2, pp. 295-312.
- Guy, M. & Tonkin, E. (2006), "Folksonomies: Tidying up Tags?", *D-Lib Magazine*, Vol. 12 No. 1, available at: <http://www.dlib.org/dlib/january06/guy/01guy.html> (accessed 1 March 2006).
- Hammond, T., Hannay, T., Lund, B. & Scott, J. (2005), "Social Bookmarking Tools (I): A General Review", *D-Lib Magazine*, Vol. 11 No. 4, available at: <http://www.dlib.org/dlib/april05/hammond/04hammond.html> (accessed 19 October 2006).
- Jörgensen, C. (2003), *Image Retrieval: Theory and Research*, USA: Scarecrow Press.

- Kipp, M.E.I. & Campbell, D.G. (2006), "Patterns and Inconsistencies in Collaborative Tagging Systems: An Examination of Tagging Practices", *Paper presented at the 2006 ASIS&T Annual Meeting, November 7, Austin, Texas*, available at: <http://dlist.sir.arizona.edu/1704/01/KippCampbellASIST.pdf>
- Macgregor, G. & McCulloch, E. (2006), "Collaborative Tagging as a Knowledge Organisation and Resource Discovery Tool", *Library Review*, Vol. 55 No. 5, pp. 291-300.
- Marlow, C., Naaman, M., Boyd, D. & Davis, M. (2006), "Position Paper, Tagging, Taxonomy, Flickr, Article, ToRead", *Proceedings of the 15th International World Wide Web Conference*, (Collaborative Web Tagging Workshop), May 22, Edinburgh, UK.
- Matthews, P. (2006), "The Fall and Rise of Collaborative Filtering", *Informr*, Vol. 20 (Autumn), pp. 2-5.
- Mathes, A. (2004), "Folksonomies – Cooperative Classification and Communication Through Shared Metadata", December 2004, available at: <http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html> (accessed 26 May 2006).
- O'Reilly, T. (2005), "What is Web 2.0: Design patterns and business models for the next generation of software", September 30th, available at: <http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html> (accessed 27 November 2006).
- Panofsky, E. (1983), "*Meaning in the Visual Arts*", Singapore: Peregrine Books.
- Quintarelli, E. (2005), "Folksonomies: power to the people", *Paper presented at the ISKO Italy-UniMIB meeting, June 24, Milan, Italy*, available at: <http://www-dimat.unipv.it/biblio/isko/doc/folksonomies.htm>
- Shatford-Layne, S. (1986), "Analyzing the Subject of a Picture: A Theoretical Approach", *Classification Quarterly*, Vol. 6 No. 33, pp. 39-62.
- Shatford-Layne, S. (1994), "Some Issues in the Indexing of Images", *Journal of the American Society for Information Science*, Vol. 45 No. 8, pp. 583-588.
- Shirky, C. (2005), "Ontology is Overrated: Categories, Links and Tags", *Clay Shirky's Writings About the Internet*, available at: http://www.shirky.com/writings/ontology_ouerrated.html (accessed 1 March 2006).
- Smith, G. (2004), "Folksonomy: Social Classification", *Atomiq: A Weblog by Gene Smith*, August 3, available at: http://atomiq.org/archives/2004/08/folksonomy_social_classification.html (accessed 8 April 2007).
- Vander Wal, T. (2005), "Explaining and Showing Broad and Narrow Folksonomies", *Personal InfoCloud.com*, February 21, available at: http://www.personalinfocloud.com/2005/02/explaining_and_html (accessed 8 April 2007).
- Wilkinson, D. (2007), "*Flickr Mashups*", Indiana: Wiley Publishing Inc.
- Wolfram, D. (2003), "*Applied Informetrics for Information Retrieval Research*", Westport, Connecticut: Libraries Unlimited.
- Wright, A. (2004), "Folksonomy", August 23, available at: <http://www.agwright.com/blog/archives/000900.html> (accessed 26 May 2006).

Appendix 1

Classification Protocol

A		Generic relationship between tag and image content
	1	Tag identifies what image is of at its most primary and objective level - no subject specific knowledge is needed to make this distinction (e.g., an image of a cat, tagged as 'cat' or 'animal').
B		Specific relationship between tag and image content
		Tag identifies what image is of . Familiarity or some existing knowledge is needed to make this connection, and to a certain extent an assumption has to be made about this connection.
	1(a)	Tags which identify place names/events – an image of the Eiffel Tower in Paris tagged as 'Eiffel Tower' requires knowledge acquired from familiarity with the specific place in question. Assumptions have to be made that an image tag is what it claims to be if the image is not familiar.
	1(b)	Tags which identify people/animals – an image of Elvis Presley tagged as 'Elvis Presley' requires knowledge and familiarity of Elvis Presley. Distinctions cannot always be made between 'famous' people and 'non-famous' people, therefore the assumption has to be made that an image of a girl, tagged as 'Sarah' is in-fact an image of a girl who is called 'Sarah'.

	2	<p>Tag identifies what the image is about</p> <p>Typically expressed by the use of abstract nouns or adjectives - an interpretation is made of what the image is about (e.g., image of people smiling tagged as 'happiness'; image of cars on a motorway tagged as 'speed').</p>
C		Tag only useful to individual/university group
	1	<p>Refining tag</p> <p>Tag which cannot stand alone - only useful when looked at as part of the larger tag set (e.g., a series of collective images, each with a specific number or letter; acronyms; dates; or information about the camera model which the image was taken with).</p>
	2	<p>Self-reference tag</p> <p>Tags which identify image content in terms of it's relation to either the tagger or the specific group which the image belongs to (e.g., 'my dog'; 'our graduation')</p> <p>OR tags which appear useful, but show no relationship/connection to the accompanying image.</p>
	3	<p>Task-organising tag</p> <p>Tags which indicate an action required on behalf of the tagger (e.g., 'move to folder...'; 'print').</p>
	4	<p>Tag which explicitly denotes ownership of image</p> <p>(e.g., image tagged with the same username as that of the person who uploaded the image).</p>
	5	<p>Compound tag</p> <p>Tags where words, phrases and sentences are joined together as one long text string.</p>
D		Miscellaneous categories
	1	<p>Misspelling (e.g., 'Belguim' instead of 'Belgium')</p> <p>Whilst it may be obvious what the tag is meant to be, a misspelling obviously renders the tag useless in terms of subsequent users of the system who are searching for images with that specific tag, unless they too misspell the tag/word.</p>
	2	<p>Unable to determine relationship</p> <p>Despite having attempted to look up either the meaning of the tag and whether the tag is a foreign word or not, tags which do not fit into any of the above categories will be deemed as unable to classify (e.g., nonsensical words).</p>
	3	Foreign word/character