Social Tagging for Digital Libraries using Formal Concept Analysis

Tim Wray and Peter Eklund

twray,peklund@uow.edu.au
School of Information Systems and Technology
University of Wollongong
Northfields Ave, Wollongong, NSW 2522, Australia.

Abstract. This paper describes the Art Collection Ecosystem – an application that allows users to tag and serendipitously browse content using Formal Concept Analysis. Within this application, tags are derived from meta-data of artworks within an existing asset management system and are classified according to theories derived from social tagging behaviour. We present past and recent iterations of its design, where it is evaluated as a contextual comparison to a popular image tagging application, Flickr. Through the process of iterative design and user evaluation, we produce results are of interest to any applied research and development that involves the exploration of digital library content using concept lattices.

1 Introduction

Digital libraries can dramatically extend the capabilities of traditional libraries by collecting, managing and preserving long term rich digital content [2]. Within the art and museum sector, projects such as Te Papa's Collections Online [1] and the Brooklyn Museum's Open Collection API [2] reflect a current industry trend of digitising their collections for public display and research purposes. There have also been some early experiments that involve the capture of social tagging data to enhance the meta-data descriptors of the collection that have been demonstrated by projects such as the steve.museum [3] and the PowerHouse Museum [4] collection databases. This approach was supported by evidence that there was a dissonance between audiences and institutions in their view, dialogue and terminology associated with art content, and that user-supplied tags could be used to fill gaps in the documentation that surround the objects [18].

Within this paper, we present the Art Collection Ecosystem, a FCA-based system that allows the synthesis and browsing of formal concepts based on a selection of artworks from the University of Wollongong's Art Collection and their

[1] http://collections.tepapa.govt.nz/
[2] http://www.brooklynmuseum.org/opencollection/api/
[3] http://www.steve.museum/
[4] http://www.powerhousemuseum.com/collection/database/

associated tags. The browsing and presentation of the art collection is based of the presentation the conceptual neighbourhood paradigm, which was presented in earlier work [7], and is detailed in section 2.2. Users have the ability to add their own tags to the works which in turn can drive and influence the underlying concept lattice, therefore influencing both collection content and navigation.

The idea of experimenting with FCA and collaborative tagging is not new. For instance, Hwang et al. describe theoretical approaches for organising and mining knowledge from user created folksonomies using FCA [12] [11]. Interestingly they have also described similar issues that we have encountered that relate the lack of semantics within the tagging process as described in [10]. Similarly, Schmitz et al. have proposed a method of using FCA as a means of mining trends and association rules within folksonomies [17]. Our work is based on the following: the design refinements and evaluations presented in section 4, particularly with respect to how the conceptual neighbourhood metaphor is used to guide users through the information space derived from tags; the application domain of this FCA-based application; and the use of a tag classification framework that allows for faceted navigation on the collection using conceptual scales.

The paper is structured as follows: in section 2, we provide a brief overview of the theory behind Formal Concept Analysis and the conceptual neighbourhood paradigm along with the tag classification framework that is used to provide faceted browsing. We then describe the Art Collection Ecosystem in section 3 and in section 4 we then follow through with two iterations of its design and user evaluation. We then conclude the paper by summarising our results and describing details of future work.

2 Background

2.1 Formal Concept Analysis

Formal Concept Analysis is a data analysis technique that allows the synthesis of formal concepts based on a collection of objects and their attributes [19]. It follows the philosophical tradition that any concept or unit of thought could be understood in terms of its attributes (or intension) and its objects that are characterised by those attributes (its extension). For instance, for a given collection of artworks, a formal concept (A, B) can be described where A represents the collection of artworks or their identifiers and B represents their set of tags (such as, { 'screenprint', 'aboriginal'}).

Formal concepts can be placed into a specialisation hierarchy, where more specific concepts (with fewer objects and more attributes) can be viewed as a specialisation of other less specific formal concepts. The result is an algebraic structure known as a concept lattice shown in Fig. 1.

Conceptual scales [9] are a powerful tool that store views of the data being analysed. Conceptual scales encompass specific attribute sets and are represented by a mathematical entity called a formal context. A context is a triple (G, M, I) where G is a set of objects, M is a set of attributes and I is an incidence relation between the objects and the attributes. Within the Art Collection

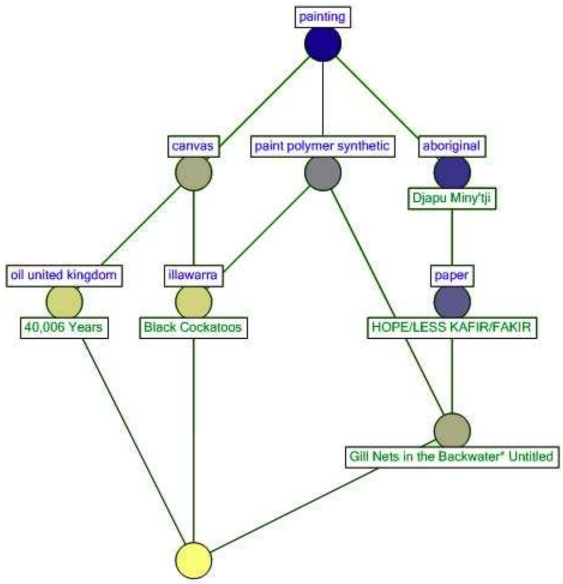

Fig. 1. A concept lattice representing a sub-context of works within the Art Collection Ecosystem

Ecosystem, we use conceptual scales to represent specific contexts with different themes such as 'artworks as described by medium' or 'artworks as described by location.' Conceptual scales can be used to combine multiple views on the data – for instance, Fig. 1 represents a sub-context created by combining the conceptual scales of 'materials' and 'location.' Within the Art Collection Ecosystem, conceptual scales are based on a classification of tag usage within collaborative tagging systems (as described in section 2.3) and users can effectively extend conceptual scales by adding their own tags to the collection.

2.2 The Conceptual Neighbourhood Paradigm

We also apply the conceptual neighbourhood paradigm for browsing the information space provided by the concept lattice. Within this approach, the user is placed at a single formal concept within the lattice. Users can move from one formal concept to another by navigating across neighbouring concepts. Fig. 2 shows a lattice neighbourhood representation of the screenshot within Fig. 3.

In our implementation of this approach, objects within a formal concept are represented as thumbnails. Users can navigate to upper neighbours (more general concepts) or lower neighbours (more specific concepts) – this is done via textual labels within the UI that are, in the case of the lower neighbours, weighted according to the extent size. Clicking on these textual labels – referred to as upper neighbour and lower neighbour controls – allow the user to move to other formal

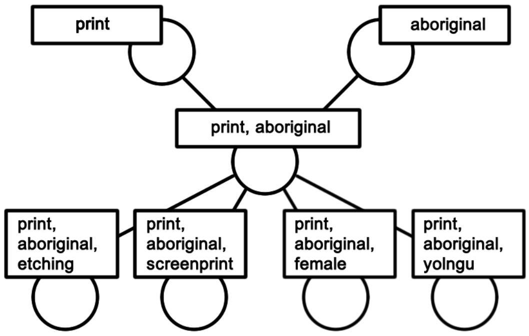

Fig. 2. The conceptual neighbourhood representation of the application screenshot as shown in Fig. 3

.

concepts within the neighbourhood. This allows the interface to move across the lattice in an intuitive way where it is impossible to navigate to a concept with an empty extent. It also allows users to change their conceptual view of the collection in minimal and incremental steps. Figs. 2 and 3 both represent a view state that describes a formal concept with attributes { 'print', 'aboriginal' } with their upper and lower neighbours. A user within this view state could easily generalise to all 'print' works or all 'aboriginal' works, or specialise to aboriginal print works that also feature keywords 'etching', 'screenprint', 'female' or 'yolngu.'

This form of navigation helps reduce the complexity at a given 'decision' point due to the natural hierarchical nature of lattice generation and navigation in which attributes can be hidden by others due to implications or attribute hierarchies. For instance, if all 'screenprint' objects are in fact a subset of all 'print' objects, then a user would be required to navigate to all 'print' objects before they could refine their query to all 'screenprint' objects, and vice-versa if they were traversing upwards through the lattice. The tagging system also supports hierarchies – if one was to tag an object as 'arnhem land' (which is a region of the northern territory) then the object will also be automatically tagged as 'northern territory' as the system stores explicit sub-type / super-type association rules between the two tags.

2.3 Classification of Tags within Collaborative Tagging Systems

Collaborative tagging systems such as Flickr have become popular in recent years due to the ability for users to annotate works with keywords that can be used for later retrieval. Tagging can be used either as a means of 'connecting' with other users and their works or as a means of organising content [1]. We focus on the latter purpose within this paper, where we use a common typology of tags based on actual tagging behaviour within collaborative tagging systems. According to Golder and Huberman [10], tags can be used to:

– Denote *subject*, that is, people, places or things being described or depicted within the resource.

- Denote *type*, whether that resource is a photo, file or bookmark.
- Denote *ownership and provenance*, identifying who owns authorship, control or rights to a resource and where it came from.
- Denote *category* or class, used to aggregate or unify resources with common properties.
- Describe specific *qualities and characteristics* of that resource, such as colour, texture, themes etc.

Within the Art Collection Ecosystem, we class conceptual scales based on these five types. For instance, the conceptual scales 'artist name' or 'materials and medium' (which contains attributes { 'canvas', 'clothing', ... }) would be classed under *qualities and characteristics*, whereas the scales of 'origin' and 'location' would be classed under *ownership and provenance*. These categories can be used to aggregate any kind of content within a collaborative tagging system. We predict that these general categorisations are useful and even necessary for managing large groups of scales or vocabularies of tags, particularly where users get to add their own tags. The classification of conceptual scales also permits a form of thematic access control: for instance, in one scenario, tags that describe an object's *ownership and provenance* are set by the museum and institution and cannot be removed or altered, or they need to rely on a fixed and standardised vocabulary to avoid duplication of terms (eg. 'Metropolitan Museum of Art', 'MMA'). However, end users are free to add their own tags that describe the *qualities and characteristics* of its content. This approach permits appropriate integration of user defined tags and content while preserving read-only (but still 'taggable') enterprise content.

3 The Art Collection Ecosystem

The Art Collection Ecosystem[5] is the latest development in a series of projects that apply Formal Concept Analysis as a means of navigating an information space. It follows from the incremental design and evaluation improvements of ImageSleuth [5], ImageSleuth2 [7], AnnotationSleuth [6] and the Virtual Museum of the Pacific [8], in that order.

Results from a previous user study of its predecessor application ImageSleuth [7] confirm the suitability of a concept lattices for the representation and navigation of image collections, and the intuitiveness of applying the conceptual neighbourhood metaphor. Although the Art Collection Ecosystem addresses some criticisms of ImageSleuth – mainly of those that relate to interface design and layout – it also employs a number of new features discussed in section 4. The objective of the user evaluations is therefore twofold: to validate (or improve) new design features of the interface but to also determine the degree of *user acceptance* – the degree to which users are likely to adopt or use a system – as a Web-based tool for browsing and exploring collection content.

[5] Available at http://epoc.cs.uow.edu.au/ace – to access the site, enter 'aceuser' as userame and 'uowace' as password.

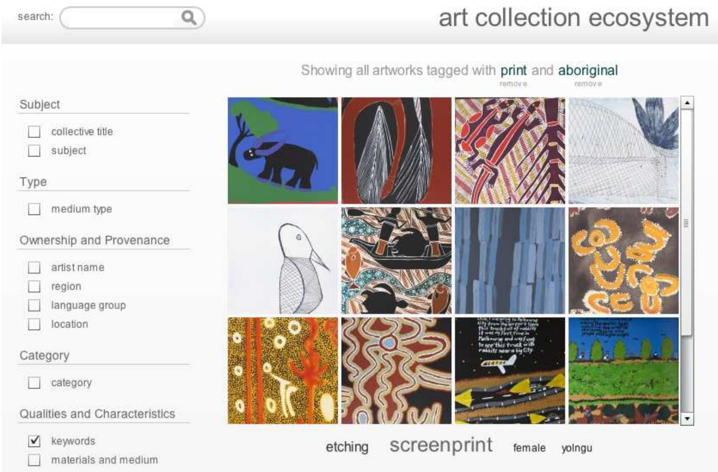

Fig. 3. A screenshot of the Art Collection Ecosystem

According to Rogers' theory of the attributes of innovations[15], a new innovation is often framed and perceived in terms of how it relates, supersede or compare to existing innovations. Therefore in our acceptibility study it is important to provide a benchmark or point of comparison where users can cast their judgements about the usefulness of an application. For this reason we will use Flickr[6] as a point of comparison: it is a modern, widely adopted image management and tagging application that encourages exploration and serendipity of its vast collections. Although we do acknowledge that the content and tasks of the two systems are different within this study, we argue that the use of Flickr as a benchmark provides users with a common point of reference on which to assess the overall design features and potential usefulness of the Art Collection Ecosystem and the concept lattice metaphor.

The study follows two iterations of a design cycle. 20 participants were used to assess and validate the original design. Some design flaws were identified and addressed. A further 5 participants were used to validate the revised design. Our conclusions were drawn from a thematic analysis of the results of the user evaluations.

4 Design and Evaluation

4.1 Initial Design

Like its predecessor [7], the Art Collection Ecosystem incorporates textual labels to provide attribute inclusion and exclusion as one would navigate to a neighbouring concept. However, it benefits from a revised design of its user interface (which was then applied to its counterpart - The Virtual Museum of the Pacific), both to improve functionality, feedback to actions and layout. These include:

[6] http://www.flickr.com

- Animation that conveys a sense of 'drilling down' and 'rolling up' as one navigates to more lower and upper neighbours respectively. To the user, this is intended to convey a sense of 'refining' and 'generalising' a particular point within an information space.
- The use of a 'tag cloud' to represent the entire collection - this is a popular means of representing key terms and their prominence.
- The hybrid combination of a tag cloud and a list of lower neighbour textual labels. Here, attribute labels are positioned and weighted according to their extent size. This allows a user to determine the prominence of a given attribute within a formal concept. In Fig. 4 for example, 'screenprint' is shown to be the most prominent attribute within the current conceptual view.
- An 'auto-complete' suggestive search field mapped to the union of all attributes within the selected scales. This was to address previous criticisms of the ImageSleuth interface regarding its ability to perform a keyword-based search.

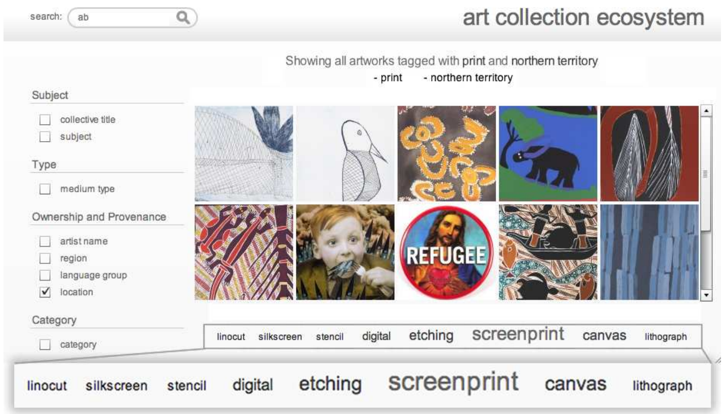

Fig. 4. Lower neighbour 'tag cloud' control within the Art Collection Ecosystem

As shown in Fig. 4 there are checkboxes on the left-hand side that allow users to toggle conceptual scales – these effectively allow users to specify and combine how they would view and browse the information space of the art collection. We class and group these scales based upon the tag classification framework discussed in section 2.3. The individual scales themselves are derived from key attribute descriptors of the artworks from the existing database of the University of Wollongong's Art Collection. As with many existing digital libraries, there were some problems of meta-data inconsistency [13], synonymy, and polysemy during the meta-data extracting process - a common problem also found in collaborative tagging systems [10].

4.2 Evaluation Cohort and Design

25 students and staff from the University of Wollongong were selected to participate in a user acceptance study. A preliminary survey was conducted to

determine the attributes of the cohort. In an answer to the question 'On a scale of 1 - 10, how much interest would you show in visiting and exploring an art collection within a web browser', the average user response was 6, although there was a high variance within this response. 18 of the 25 surveyed users have used or were already familiar with Web-based tagging applications. The use of Flickr as a point of comparison was intended to standardise the level of experience users had with existing Web-based image exploration and tagging applications, along with providing a baseline point of comparison as described in section 3.

Users were required to perform tasks that involved searching and exploration in both applications. Screencams and commentary were recorded as users performed these tasks. A qualitative analysis was done where recurring themes and issues were identified. 20 users undertook an initial study, then a further 5 users were surveyed to validate the design changes described in section 4.4. After the evaluation sessions, all users were required to complete a brief survey that documented their overall impressions with the Art Collection Ecosystem.

4.3 Initial Evaluation Results

After the initial evaluation, users generally held a positive attitude towards the design and functionality of the user interface for the Art Collection Ecosystem. In comparison to the functionality afforded by Flickr, users generally praised both the aesthetics of the user interface along with the ability to 'hold' a current or particular point or concept within the information space. This finding is consistent with the previous usability study on ImageSleuth [7]. The ability to generalise or refine's one view within the information space contrasted with the search functionality offered by Flickr where users had trouble refining their search terms – for instance, if they wanted to select a particular sub-set of images within their results, users were required to alter their text string within the search field, often generating a result set that was disjointed from the original.

Users generally praised (and understood) the effect and usefulness of having the ability to toggle conceptual scales in order to define a context. More specifically it allowed them to specify the dimensionality of their search – such as whether they wanted to search by material type, artist name, or physical location. Many perceived these controls "as a way of refining your search" and made comparisons with this and Flickr's 'Advanced Search' functionality. However, the implementation and categorisation of conceptual scales into groupings such as 'Subject', 'Type' and 'Ownership and Provenance' was met with a mixed response. There were positive remarks with respect to clear demarcations of scale options, although there was some criticism with respect to ambiguity to the names of the scales (for instance, there was a scale called 'location' and another called 'region') and there were some comments on the triviality of the feature. Further criticism was directed at the meaning behind individual tags and attributes – in expressing concerns of tag polysemy, one user noted that the tag 'united kingdom', for example, could refer to the current location of the work, where it was sourced from, or the birthplace of the artist. This problem is caused by the lack of semantic information in the tagging process [14] even though the purpose of using conceptual scales was to seperate and filter the varying types of tag-to-resource associations.

In terms of conceptual neighbourhood navigation, 18 out of 20 users immediately understood the purpose and effect of the lower neighbour control (as shown in Fig. 4), although many users thought that the size of the tag denoted popularity (ie. the amount of times it was clicked), rather than the neighbouring concept's extent size. However, only 6 out of the 20 users correctly interpreted the purpose and the function of the upper neighbour control. In an example shown in Fig. 5, the remaining 14 users thought that the effect of the control was that they would simply navigate to all objects tagged 'northern territory' and 'print' rather than navigate to an upper neighbour that excluded that particular attribute. Additionally, 4 out of the 20 surveyed users also expressed the need for a 'Back' button. This is due to the hierarchical but non-linear mode of navigating the information space. These users sometimes felt 'lost' and expressed a need to 'go back' and traverse a list of previously navigated conceptual views.

4.4 Design Changes

The initial results of the study have motivated two design changes to improve interface learnability and functionality.

The first and most pressing of these of these changes was the revision of the upper neighbours control. Figs. 5 and 6 show two versions of the upper neighbour control displaying the same navigation links.[7] As shown in these examples, users are given the opportunity to navigate to an upper neighbour by excluding attributes 'print' or 'northern territory' from their current conceptual view.

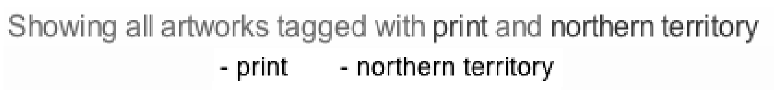

Fig. 5. Upper neighbour control, before design revision

Fig. 6. Upper neighbour control, after design revision

The revised control demonstrates a more natural means of doing this. In Fig. 6, the description of a formal concept's intension contains emboldened hyperlinks that highlight which attributes are removable. As a user rolls their mouse over the link, a strikethrough effect hints the intended action of clicking on the link.

The second design change involves the placement of a 'Back' button, allowing users to traverse their conceptual navigation history, which is a list of view states.

[7] Please note that the colour of the navigation labels have been altered for print clarity

Each view state contains: a) a reference to a formal concept as described by its intension and b) the current 'activated' scales that were selected by the user. This is so that a user can consistently navigate backwards and forwards through their view history, even if the formal context changes as a result of tagging activity.

4.5 Secondary Evaluation Results

Some of the findings and common themes of these results were consistent from the first study – these included the positive aspects of interface aesthetics along with the ability to 'hold' or retain a particular position within the information space.

All 5 users recognised the function and purpose of the upper neighbours control almost immediately, and none alluded to the incorrect assumption described in section 4.3. Results concerning the 'Back' button however, were less prominent – only 2 of the 5 users saw a need to use it. However, both of these users said it was necessary to have the ability to traverse backwards through the information space.

5 Discussion of Results and Future Work

The results of the study indicate a moderate level of user acceptance of the use of the Art Collection Ecosystem in browsing and exploring tagged content. The initial around of 20 users provided an average score of 0.3 in response to the question "I would prefer to use the Art Collection Ecosystem to explore objects than I would with Flickr", were a score of -1 denotes 'strongly disagree' and a score of 1 denotes 'strongly agree'. Users generally associated the conceptual navigation paradigm as a means of query refinement and content exploration, where an average score of 0.9 was given for the question "It was easier to browse serendipitously or find related images within the Art Collection Ecosystem than it was in Flickr." Subjective and aesthetic commentary was also given, with many users claiming that the user experience of conceptual navigation was more "hands on" and "involved." Users also noted that, compared to Flickr, it may require some time and experimentation to understand conceptual neighbourhood paradigm. Furthermore, the validated design features of the Art Collection Ecosystem have also been implemented in another software application that shares the same framework - The Virtual Museum of the Pacific [8].

One of the major themes of the results centered on the learnability of the two interfaces – a major facet that determines a user's willingness to adopt or use an application [16]. Users commented that although the Art Collection Ecosystem presents itself as a useful means for query construction and refinement, many would state that they would be willing to a more familiar tools that facilitate search and retrieval, such as the use of a search box, unless they had a good reason to invest the time required to become familiar with the user interface through experimentation.

Even though our *initial* meta-data source for our case study was from a single (non-public) repository of data, the sometimes inconsistent usage of terminology – perhaps as a result of a shift between departments or individuals who maintain it – can affect the quality of the browsing experience. This is a classic problem of knowledge management [3] within many digital libraries [13]. These issues of meta-data noise and inconsistency can strongly affect the quality of results within a search query and can be severely detrimental to the quality of results in an application that employs concept lattices for search and navigation [4]. We argue that conceptual scales can be used to classify tags according to usage within collaborative tagging systems (as discussed in sections 2.1 and 2.3) – this in turn allows users to toggle specific dimensions of their navigation context or filter out undesirable noise. We postulate that the same approach could be used to toggle and separate between well defined and consistent enterprise data apart from user contributed (but sometimes noisy or inconsistent) tags and attributes. Future work and longitudinal studies based on tagging behaviour will investigate problems of navigating digital libraries using noisy non-mediated vocabularies.

Future work will also assess the scalability of our approach, with possible methods of using horizontal scaling to cope with the conceptual neighbourhood computation of increasingly large data sets and contexts. This is also necessary to incorporate the expected exponential growth of custom objects and user-defined tags if users were to adopt this as a content presentation and tagging application. We are currently planning to expand the data source of the Art Collection Ecosystem with approximately 2500 artworks, representing a collection complete implementation of an FCA-based application.

6 Conclusion

In this paper we have presented the design and evaluation results of the Art Collection Ecosystem. The results of our study reinforce the viability of using the conceptual neighbourhood approach in browsing tagged image collections. The results have also driven and validated design improvements. A useful application of applying conceptual scaling to segment tag vocabularies based on a framework of collaborative tagging usage was also discussed, driving forward the development of a collaborative FCA-based digital library application.

References

1. D. Boyd, M. Davis, C. Marlow, and M. Naaman. HT06, Tagging Paper, Taxonomy, Flickr, Academic Article, To Read. In *Proceedings of the ACM Conference on Hypertext and Multimedia*, Odense, Denmark, 2006.
2. L. Candela, D. Castelli, N. Ferro, Y. Ioannidis, G. Koutrika, C. Meghini, P. Pagano, S. Ross, D. Soergel, M. Agosti, M. Dobreva, V. Katifori, and H. Schuldt. The delos digital library reference model - foundations for digital libraries. Technical report, DELOS Network of Excellence on Digital Libraries, 2007.

3. T. Davenport and L. Prusak. Information politics: Information is not innocent. In *Information Ecology: Mastering the Information and Knowledge Environment*, chapter 5, pages 67 – 82. Oxford University Press, New York, 1997.
4. J.. Ducrou. *DVDSleuth: A Case Study in Applied Formal Concept Analysis for Navigating Web Catalogs*, pages 496–500. LNCS 4604. Springer-Verlag, 2007.
5. J. Ducrou, B. Vormbrock, and P. Eklund. FCA-based Browsing and Searching of a Collection of Images. In *Proceedings of 14th International Conference on Conceptual Structures*, LNAI 4068, pages 203–214. Springer, 2006.
6. P. Eklund and J. Ducrou. Navigation and annotation with formal concept analysis. In *Knowledge Acquisition: Approaches, Algorithms and Applications*, LNCS 5465, pages 118 – 121. Springer Berlin / Heidelberg, 2009.
7. P. Eklund, J. Ducrou, and T. Wilson. An intelligent user interface for browsing and search MPEG-7 images using concept lattices. In *Proceedings of the 4th international conference on concept lattices and their applications*, LNCS 4923, pages 1–22. Springer-Verlag, 2008.
8. P. Eklund, T. Wray, and J. Ducrou. Web services and digital ecosystem support using formal concept analysis. In N. Spyratos, editor, *Proceedings of the International Conference on Management of Emergent Digital EcoSystems*, pages 238–245. ACM Press, 2009.
9. B. Ganter and R. Wille. Conceptual scaling. In F. Roberts, editor, *Applications of Combinatorics and Graph Theory to the Biological and Social Sciences*, pages 139–167. Springer-Verlag, 1989.
10. S. Golder and B. Huberman. The structure of collaborative tagging systems. *Journal of Information Science*, 32(2):198 – 208, 2006.
11. Y. Kang, S. H. Hwang, and K. M. Yang. FCA-based conceptual knowledge discovery in Folksonomy. *World Academy of Science, Engineering and Technology*, 53:842 – 846, 2009.
12. H. L. Kim, S. H. Hwang, and H. G. Kim. FCA-based Approach for Mining Contextualized Folksonomy. In *Proceedings of the 2007 ACM Symposium on Applied Computing*, 2007.
13. C. Lagoze, D. B. Krafft, S. Payette, and S. Jesuroga. What is a digital library anymore, anyway? *D-Lib Magazine*, 11(11), 2005.
14. A. Marchetti, M. Tesconi, and F. Ronzano. Semkey: A Semantic Collaborative Tagging System. In *Proceedings of the WWW Workshop on Tagging and Metadata for Social Information Organisation*, 2007.
15. E. M. Rogers. *Diffusion of Innovations*. The Free Press, New York, NY, USA, 1995.
16. Evelyn P. Rozanski and Anne R. Haake. The many facets of HCI. In *CITC4 '03: Proceedings of the 4th conference on Information technology curriculum*, pages 180–185, New York, NY, USA, 2003. ACM Press.
17. Christoph Schmitz, Andreas Hotho, Robert Jäschke, and Gerd Stumme. Mining association rules in folksonomies. In Vladimir Batagelj, Hans-Hermann Bock, Anuška Ferligoj, and Aleš Žiberna, editors, *Data Science and Classification*, Studies in Classification, Data Analysis, and Knowledge Organization, pages 261 – 270. Springer Berlin Heidelberg, 2006.
18. J. Trant and B. Wyman. Investigating social tagging and folksonomy in art museums with steve.museum. In *Proceedings of the International World Wide Web Conference*, Edinburgh, Scotland, 2006.
19. R. Wille and B. Ganter. *Formal Concept Analysis: Mathematical Foundations*. Springer-Verlag, Berlin, 1999.