

On the Importance of “Who Tagged What”

Mark van Setten, Rogier Brussee, Harry van Vliet, Luit Gazendam,
Ynze van Houten, Mettina Veenstra

P.O. Box 589, 7500 AN, Enschede, The Netherlands
{Mark.vanSetten, Rogier.Brussee, Harry.vanVliet,
Luit.Gazendam, Ynze.vanHouten, Mettina.Veenstra}@telin.nl

Abstract. Adaptive systems can be used to help people find their way through large amounts of information by taking into account the content itself and opinions of other people. Annotations that describe content are important in this process as is generally recognized. What is less obviously stated so far is that not only the annotations themselves, but also knowledge of who provided the annotations can be relevant for adaptive systems. Annotations can namely be regarded as an opinion of the annotator, which can only be valued if one knows who provided it. Various sources of annotations and the importance of knowing who provided an annotation are described using real-world cases. This is followed by the introduction of various ways that adaptive systems can use the knowledge of “who tagged what” when helping people find their way through large amounts of information.

1 Introduction

One of the main concerns of adaptive hypermedia is to support people in finding their way through large amounts of (linked) content by automatically selecting, structuring or suggesting content while taking into account the user’s knowledge and/or preferences [2]. Such adaptive information systems focus on content that the system deems interesting for the user thereby relieving some strain from the user’s limited attention span. The decision which content is interesting and which is not is generally heuristically determined by taking into account metadata of the content and opinions of users. This is the area most research on recommender systems focuses on [9].

We argue here that annotations added to content can also be regarded as an opinion. This way, in addition to the annotations themselves, the relation of the user towards the annotator provides valuable information to determine whether content is interesting for a user or not. The basis for this argument is the idea that annotations are opinions or views of the annotator rather than objective facts. Such an opinion or view is more valuable if one knows who provided it, because it allows a better founded assessment of the annotations or opinions.

In this paper, we address the importance of knowing who provided annotations by using several cases from projects in which the authors of this paper participated; these cases are presented in section 3. In section 4, the implications and possibilities of

keeping track of who provided annotations for adaptive systems are discussed. We start the discussion in section 2 by describing several types of annotations.

2 Types of Annotations

Annotating is the process of adding additional information to a document or other electronic objects. The traditional annotation is a comment in the margin. In the digital domain, annotations can range from bibliographical metadata such as author, title and date of publication to technical metadata such as file format, and from user defined keywords to computer generated abstracts. In this paper, we focus on annotations that are intended to facilitate the retrieval of documents and other electronic objects.

We distinguish four different types of annotation sources: annotations made by the creator, annotations provided by cataloguers, social annotations and annotations provided by systems.

2.1 Annotations Made by the Creator

Creators of content often provide annotations as part of their regular creation process, even though it can be time consuming. The motivation to provide them is to increase the chance that their content will be found by others. An author of a book for example, determines the title and (usually) puts his name down as the author. Likewise, the table of contents can be regarded as metadata about the content of a publication. Publishers, who are also part of the publication process, also add metadata to content, such as identification numbers, publication dates and summaries. Apart from the last example, such metadata is objective or at least, by convention or law, the responsibility of the publisher. Other annotations, like reviews printed on book covers, may contain a subjective positive bias towards the publication. In printed media such a bias is harmless, because the conventional form of such reviews makes it apparent for a possible buyer how to assess such a cover text, e.g. by taking into account that quotes are selected by the publisher, and the source of such a review.

2.2 Annotations Provided by Cataloguers

Cataloguers are professionals whose job it is to store documents in archives for easy retrieval in the future. One of the traditional ways of archiving is by categorizing objects on subject. Such annotations are widely used, and classification schemes designed for this purpose, such as the library of congress subject classification or the Dewey decimal system, have a long tradition in library practice and science. The attraction of a formal classification system or a controlled and often structured vocabulary (a so-called thesaurus [1]) lies in the great simplification of the categorization process. Other documents with exactly the same keyword are naturally grouped together. Likewise, documents assigned a hierarchical category are themselves grouped in a hierarchical fashion. Such a grouping is very useful to find

related content, and it is trivial for a computer to find documents that have exactly the same keywords or fall into an exactly defined category. However, it does not address the following *non technical* problems:

1. How to ensure that content is *properly* annotated with keywords and/or categories?
2. How to ensure that those keywords or categories reflect what a user is looking for?
3. How to decide on the classification scheme or thesaurus?
4. How to know that a keyword or category was assigned in the intended or at least in a reasonable way?

Addressing all these issues by cataloguers alone is expensive (especially with the continuously increasing amount of accessible content), difficult to achieve and requires detailed knowledge of users. Other sources of annotations have therefore acquired attention. In addition, for textual content, full text search is used routinely for tasks that otherwise would have required very detailed annotations. Full text search is user driven and retrieval is based on what authors have written, sidestepping the above problems. The price to pay is that it is hard to judge how a query is reflected in the search result and people have to learn how to best formulate a query that will lead to their higher order goals. Fortunately, many people have become experienced with search engines and are aware of the pitfalls. However, content-based retrieval is still extremely difficult for non-textual media (see section 2.4).

2.3 Social Annotations

Social annotation is another way of acquiring annotations. The idea is to provide an open system that elicits people to annotate or tag people’s content, especially their own content, often resulting in so-called folksonomies [7]. Such a process is greatly facilitated by easy access, availability, standardization at the same level of annotations and attractiveness of the medium. Openness means that, in principle, there is no hierarchy or initial weighing in who contributes annotations and there is no real barrier to contribute information (‘everyone is equal, everyone can tag’). In practice this openness can be restricted by the author to a peer group, a maintainer or the author himself. There may be value in sharing the content and its annotations to a group, but that does not imply that it has to be everybody on the Internet. As getting oneself noticed and interacting with others appears to be a reward for active annotators, open annotation systems tend to become communication systems with social interactions as a consequence, see for example www.flickr.com.

Social annotation is in some way a misleading term because there is nothing inherently social about social annotation. It has consequences which can be deemed social; that is, there are consequences related to this loosely coordinated behavior of people. For instance, openness means many people join in and many hands make light work. Openness also creates opportunities for new social interactions, like finding others with common interests. If we consider these as positive impacts of openness, there are also negative impacts, like vandalism. Every open system is vulnerable to misuse by individuals, which can easily disturb an open system; not only making it unstable but also changing the way it is experienced, i.e. as non reliable. For instance, rewarding people for their participation in online communities by way of the quantity of their contributions leads to people sharing resources of poor quality or unrelated to

the topic, resulting in a tougher job to find good resources and in the end to a decreased level of participation [3]. In many systems, there are processes at work to ‘neutralize’ such vandalism. For example, the Wikipedia relies on guidelines for well-meaning contributors to behave responsibly and write as accurately and objectively as possible, the possibility of immediate feedback to ‘correct’ pages, and a policy to return pages to their original state as soon as they are recognized as vandalized [12]. However, such maintenance can become a burden. Another side effect is ‘dominance’: some annotations will climb to the top of most used tags because a vocal user community will tag a disproportionate amount of content or because of their ‘track record’ always come up on top resulting in ‘the rich get richer’, also called the Matthew effect [3]. This rise in tags then starts to bias the computed relevance of content, i.e. their ‘bias’ resonates through the system. It can also result in people reusing other people’s tags without being aware of the context in which the original users provided those tags; del.icio.us even suggests other user’s tags as tags to use for annotating websites. This can heavily influence the appraisal of what content we find (relevant). In a study of the Slashdot rating mechanism (slashdot.org) it was shown that comments have initial lower ratings and those contributed late in the discussion receive insufficient attention and end up with an unfair score, despite deserving comments [3]. Finally, the ‘democracy’ of the openness can also lead to possibly unwanted annotations. For instance we do not expect a ‘voting’ on the authors name or the title of a published book, whereas such a voting is expecting on issues such as whether it is ‘exciting’ or ‘the best book of the year’. Such ‘factual’ annotations as title are in some way different from ‘opinion’ annotations, and with such factual information there are certain people who are the authentic source of it (the author, the publishing company). So some annotations must be protected from popular belief (the majority), who might think otherwise for some reason than the authentic source of information (the minority).

2.4 Annotations Provided by Systems

Systems that automatically analyze and annotate content without human supervision are an active area of research. Unsupervised metadata creation uses algorithms to automatically detect content characteristics. While being attractive because it requires limited human effort, it runs into the problem of the “semantic gap” (especially for non-textual content): the problem that semantic interpretation of content requires more world knowledge or even computer power than computers currently have. In the visual domain, this is particularly acute as annotations are often the only handle to retrieval due to the lack of correspondence between the information extractable from the visual data and the interpretation that the data has for a user in a given situation [10]. Clearly, higher-level semantic descriptions are often more useful than low-level features like color and texture, but automatic classifiers for such high-level features are much less accurate [8]. Promising developments include multimodal analysis techniques, using data from the visual, auditory and textual modality [10]. Information extraction techniques for text [4] are better developed but still run into lack of context.

3 Case Studies

In this section, we discuss four cases in where various sources of annotations were used by content providers for a variety of reasons. The authors participate(d) in projects within those cases on a variety of research topics including ways to support professional cataloguers in the annotation process, automatic classification techniques, video browsing, automatic annotation techniques, design studies, social annotation techniques, etc. The viewpoints as described in the cases are all based on either results of those research projects or based on interviews and discussions with members of the organizations or communities in which the cases are based.

3.1 Ondergrond.org

In the fall of 2005, the Telematica Instituut gave an assignment to students of the European Media Master of Arts curriculum of the Utrecht School of the Arts to visualize how people can participate in supporting our cultural heritage. They decided to create a website called www.undergrond.org devoted to street art in public places, such as graffiti, posters and stickers. Street art is typically spread across a city, tied to a specific location. It is also subject to weathering, cleaning or being sprayed over, making it a fleeting media that leaves few traces. What started as an experiment has quickly grown into a website with several features. It provides the possibility to upload photos, requiring a minimal annotation of the object such as location, a short description and some keywords. It can also be extended with, amongst other things, artist, crew, color and medium. This allows the creation of different ‘expositions’, for instance based on location, artist, color or keywords (humor, politics, etc). Such expositions are constantly present in the interface, as a way to lure visitors into an associative chain of clicks through the overall collection.

To determine and increase the quality of tags, a Q&A-module is provided asking visitors to confirm or vote on certain metadata. For instance, a question could be asked whether the color of the object is predominantly red. People may also be asked whether a keyword seems to best describe the object with some choice provided from keywords added by others. The underlying Q&A system is build around certain categories of questions and monitors if all different questions are evenly addressed.

Ondergrond.org relies almost exclusively on social annotations, reflecting the community-style of street arts. Both artists and interested others can (and do) provide annotations and actively participate in the community. Although the system registers who is providing information and tags it currently does not show who provided which annotations. According to feedback from the community, which actually consists of several sub-cultures with their own fiercely uphold style, views and opinions, this would be a very welcome addition.

3.2 Fabchannel

Fabchannel (www.fabchannel.com) broadcasts concerts, festivals, competitions and lectures on the Internet from two Amsterdam concert halls (Paradiso and the

Melkweg). They also offer one of the largest live music streaming video archives with over 400 recordings. Apart from big international acts, Fabchannel is always on the lookout for new, less known, but talented bands on the edge of breakthrough and helps them to reach a broader audience by spreading their music.

The Fabplayer is a web-based video player framework developed by Fabchannel's R&D department designed to explore the archive. It supports intelligent video browsing within recorded concerts. Current developments are aimed at supporting users to efficiently navigate through video content with the aid of annotated segments grouped on the basis of similar attributes such as guitar solos or all songs from the same CD, a form of interaction called patch-based video browsing [6].

The first step towards obtaining annotated videos is their segmentation. Currently, concerts are manually segmented into songs, but work is in progress to automate this process. Development of algorithms for the detection of specific song parts like instrument solos, applauses and 'exciting' segments is also planned. Song segments can be linked to external sources, such as music databases with song and artist related data (album, composer, year, etc) or commercial music download sites. Currently, professional cataloguers of FabChannel provide the annotations for song segments based on information such as set lists. However, as sets lists are not always accurate and the cataloguers are generally not familiar with the bands, it requires a lot of research from the cataloguers to annotate the song segments accurately. Based on interviews with members of FabChannel, we know that for this reason FabChannel intends to invite band members to annotate their own performances and to allow users of FabChannel to provide annotations and to correct existing annotations. These users are fans from all over the world, and fans generally know a lot about the artists and their performances.

3.3 Sound and Vision

The Netherlands Institute for Sound and Vision (www.beeldengeluid.nl) archives Dutch public radio and television programs for their testimonial value (in the cultural-heritage sense) and for the reuse by program makers. Since March 2006, all programs created in the Netherlands and broadcasted by the public broadcasting companies are automatically submitted to their digital archive. Sound and Vision makes a distinction between high and low priority programs. Low priority programs such as game shows, entertainment programs and comedies are currently only annotated on a basic level using metadata such as title, broadcasting date and genre generally provided by the broadcasters. High priority programs such as news broadcasts, current-affairs programs and documentaries, are annotated in more detail by professional cataloguers.

During the annotation process, cataloguers consult a set of textual resources to identify the most relevant categories for a given program and when necessary watch (parts of) the program. Categories are chosen from a thesaurus developed over the years by Sound and Vision. The available resources for having cataloguers annotate programs are limited; there are many high priority programs and a fixed number of cataloguers. Due to decreasing government budgets and an increasing number of programs that need archiving, Sound and Vision is currently investigating automatic

metadata extraction techniques; fully automated for low priority programs and semi-automated for high priority programs (semi-automated means suggesting annotation terms to cataloguers by automatically analyzing textual context documents).

Program makers that look for program segments to reuse are Sound and Vision’s main user group. It is important for them that programs are annotated in a consistent and reliable way; hence the use of professional cataloguers. However, in the future Sound and Vision intends to open up their archive to other user groups, including the research community and the general public. Their annotation needs are probably different. For this reason, Sound and Vision is likely to further explore other sources of annotations, such as automatic annotations and probably social annotations.

3.4 Cleco

The Cleco project explored the automatic annotation of educational content. The project was part of Kennisnet’s “Content keten (chain)” program to stimulate the creation of high quality educational content (see www.kennisnet.nl). The e-learning community puts great emphasis on annotation and uses compound document formats (learning objects) that put metadata on almost every level. As a result, the annotation process is either expensive or does not get done. Both publishers and authors are therefore anxious to get support from automatic metadata generation. A workshop with a wide variety of participants from the educational domain showed that many participants think that an automatic solution is more attractive than annotations by non-experts and they had great expectations of the possibilities of automatic metadata generation.

The Cleco system automatically extracts Dublin core type metadata from textually oriented content: language, keywords, title, author and a summarization. A metadata editor allows human alterations and/or additions to the metadata, making sure that the final set of metadata clearly indicates the sources of annotation. First responses to the results were mixed, mainly due to the lack of available algorithms that focused on the Dutch language. Several participants at a workshop, both publishers and instructors, expressed interest in being able to see the difference between human and computer generated annotations. One participant noted that this would be possible by using several metadata documents in a compound document authored by different people or computer systems, and that this was an existing best practice for IEEE-lom (<http://ltsc.ieee.org/wg12>) metadata. However, it is unclear if this distinction is and will be supported by the available learning content management systems in practice.

4 Annotations and Adaptation

In the cases presented, annotations are necessary when searching for information, where the value of annotations depends on who is searching and what the goal of that person is, i.e. annotations are contextualized. Adaptive systems are designed to take into account the context of its users, the user’s background and the user’s goals, and to provide a service that automatically adapts to the user’s context. Here, we will consider several ways in which information systems can adapt using annotations and

thus better support users in finding information to meet their needs. However, one basic requirement is that the context of an annotation is stored with the annotation: who provided the annotation (the person and his background) and the intent of the person who provided the annotation.

One possibility for an adaptive system to adapt to the context of annotations is to track which annotation contexts a user is most interested in or which annotation contexts the user trusts most. Those annotation contexts are then used exclusively or primarily when searching for content by that user. Determining which annotation contexts to use for a specific user can be done either by the user himself or by letting the information system learn which annotation contexts suits that user. When users specify their own preferred annotation context, stereotype descriptions could be used. E.g. this would allow broadcasting employees to inform the search engine of the Sound and Vision archive to primarily use the ‘professional’ stereotype which focuses on annotations provided by the professional cataloguers. Fully adaptive systems are able to learn which annotation contexts are best suited for each user. Those systems analyze the user’s past behavior to determine which annotation context best suit that user (perhaps even taking into account the user’s current goals). Based on this past behavior, an adaptive system can decide to only use those annotation contexts that it believes suits the user best or it can weigh the available annotation contexts based on the probability of fitting the current user.

In information systems where people who search for content also provide annotation themselves, adaptability can be achieved by using annotations as an opinion of the annotator. Based on this idea, if someone takes the trouble of making annotations himself, those annotations can become part of his user profile as an indication of his perspective on the content collection and interest in the annotated object. This profile can then be used in a way similar to how rating profiles are used in recommender systems. A recommendation technique such as collaborative filtering [5] can be used to find users who annotated objects similarly. These similar users can then be used to suggest content that has been annotated by those users but not by the current user. It can even be used to suggest annotations to a person for objects he has not annotated before or additional annotations to objects he has already annotated.

Similarly, other recommendation techniques can also be used for annotations, such as the case-based reasoning recommendation technique [9]. Case-based reasoning finds objects that are similar to objects the user has seen in the past and uses the user’s opinions about those similar objects to determine the expected opinion of the user for the new objects. When applied to annotations, case-based reasoning can suggest possible annotation terms to the user, based on how it annotated similar objects in the past. Similarity of objects can be determined by using any combination of already available annotations. Another use would be to, instead of wait till the user provided annotations, assume that the annotations with the highest probability suggested by these recommendation techniques are from the user and use them when the user searches for information.

Another possibility would be to take the user’s opinion of an annotator into account. For example, if someone uses or approves of content that has been annotated by someone else (even if that user would annotate the object differently) this indicates that the annotations of that person have some value to the user; i.e. the rating for or usage of an object propagates as a rating on all users who provided annotations for

that object. This is an implicit way of endorsing other users. Endorsement can also be done explicitly by allowing users to give a rating to or opinion on other users.

Adaptive systems in combination with annotations offer a lot of possibilities to help users in their use and search for information, either by suggesting (additional) annotations that a user might want to add to an object or by determining which annotation contexts can best be used during a search for information and thus improve the quality of search results.

5 Conclusions

In this paper, we have argued for the importance of knowing who provided which annotations and gave anecdotal evidence that this is a problem of interest to a user group based on several real-world case studies. We distinguished four different types of annotations based on the roles of people or systems that were responsible for their creation. In the cases presented, we discovered that all four types of annotations can provide useful annotations for certain groups of users and stakeholders. Using one source predominantly does not mean that other sources of annotation are useless; in fact, it seems likely that different types of annotations and multiple annotations by different people or systems will find simultaneous use. This increases the need to make the source of annotations explicit. For some types of annotations, like ratings, or the result of a search engine, people already expect that annotations do not give a clear cut, authoritative answer. Thus the challenge is to deal with ambiguous, incomplete, contradictory, subjective and partly comparable annotations from different sources. This situation seems to be just a fact of life. Taking the source of information into account is often used as a coping strategy when information is ambiguous, incomplete or subjective. Thus we believe that knowing where annotations come from also helps to improve judgment.

Knowing who (or what) provided the annotations can help people in determining the relevance of the annotations for their own goals. Often this means that users have to have additional information on the annotator, e.g. role and affiliation, human or machine, other annotations or work of that annotator etc. What will be useful characteristics for annotators clearly depends on the context: .e.g. in few cases will it matter whether an annotator spends nights spraying graffiti, but in the *ondergrond.org* case it clearly does. For the professional user community at *Sound and Vision*, it is important that programs such as news broadcasts have been annotated by experienced professional cataloguers in a standardized way (in this case using the thesaurus) as this provides uniformity in annotations and a reliable level of quality. If TV programs would also be annotated by other annotation mechanisms, e.g. automated annotation techniques or social annotation techniques, the professional cataloguers and users should be aware of this. Users of e-learning content expect publishers to provide very detailed annotations, but publishers find this extremely expensive. Automatic annotation can satisfy their user’s needs to a certain extent, but both users and publishers expressed interest in making it clear that those annotations are computer generated. Even for people searching through concert videos and people browsing through the street art collection, knowing who provided certain annotations can help

them better understand the meaning and thus value of the annotation for their own purposes as some of the annotations are arguably subjective

As adaptive systems are meant to provide services that are tailored to each individual user and his goals, thus taking care of subjective perspectives, we introduced several possibilities of how adaptive systems can use annotations and the knowledge of who provided those annotations. Such adaptive systems will be able to better help people find their way through large content collections. Our intention is to investigate these possibilities in future research.

Acknowledgements

The authors like to thank the students of the Utrecht School of the Arts, Fabchannel, The Netherlands Institute for Sound and Vision and Kennisnet for their participation. We also thank Christian Wartena for his useful comments.

References

1. Aitchison, J., Gilchrist, A., Bawden, D.: *Thesaurus Construction and Use: A Practical Manual*. Fitzroy Dearborn Publishers, 4th edition (2002)
2. Brusilovsky, P. Methods and techniques of adaptive hypermedia. *User Modelling and User-Adaptive Interaction* 6, 2-3 (1995) 87-192
3. Cheng, R., Vassileva, J.: Adaptive Reward Mechanism for Sustainable Online Learning Community. In: *Proceedings AI in Education, AIED 2005, Amsterdam, 18-22 July 2005* (2005) 152-159
4. Cunningham, H.: Information Extraction. *Encyclopedia of Language and Linguistics* (2005). Available at <http://gate.ac.uk/sale/ell2/ie/main.pdf>
5. Herlocker, J., Konstan, J.A.: Content-Independent Task-Focused Recommendation. *IEEE Internet Computing* 5 (2001) 40-47
6. van Houten, Y., Schuurman, J. G., Verhagen, P.: Video content foraging. In: Enser, P., Kompatsiaris, Y., O'Connor, N., Smeaton, A.F., Smeulders, A.W.M. (eds.): *Lecture Notes in Computer Science 3115 - Image and Video Retrieval - Proceedings of CIVR 200*. Springer-Verlag (2004) 15-23
7. Mathes, A.: Folksonomies – Cooperative Classification and Communication Through Shared Metadata. Online: <http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html>
8. Sebe, N., Lew, M. S., Smeulders, A.W.M.: Video retrieval and summarization (Editorial introduction). *Computer Vision and Image Understanding* 92 (2003) 141-146
9. van Setten, M.: *Supporting People in Finding Information: Goal-based Structuring and Hybrid Recommender Systems*. Telematica Instituut, the Netherlands. ISBN 90-75176-89-9
10. Smeulders, A.W.M., Worring, M., Santini, S., Gupta, A., Jain, R.: Content-based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22 (2000) 1349-1380
11. Snoek, C. G. M., Worring, M.: Multimodal video indexing: a review of the state-of-the-art. *Multimedia Tools and Applications* 25 (2005) 5-35
12. Wikipedia: Dealing with Vandalism. Online: http://en.wikipedia.org/wiki/Wikipedia:Dealing_with_vandalism