2009

Using social networking and collections to enable video semantics acquisition

S. J. Davis
University of Wollongong, stdavis@uow.edu.au

Christian H. Ritz
University of Wollongong, critz@uow.edu.au

I. Burnett
Faculty of Informatics, University of Wollongong, ianb@uow.edu.au

Using social networking and collections to enable video semantics acquisition

Abstract

Video content both professionally and user generated is now a prime driver of Internet usage and activity; Simultaneously, the Internet has seen the rise of social networking sites and activity. This paper brings together video content and social networking to solve the problem of gathering metadata which describes user interaction, usage and opinion of video content. The paper explores mechanisms by which implied and explicit semantics can be garnered from users interacting with media within social networks. This leads to the gathering of individual user interaction metadata which is then aggregated to form semantic metadata for a given video. The techniques have been implemented in a custom Flex application which is based around the Facebook API. This provides various mechanisms for the collection and then display of descriptions of user interaction with video content.

Disciplines

Physical Sciences and Mathematics

Publication Details

Using Social Networking and Collections to Enable Video Semantics Acquisition

Stephen J. Davis and Christian H. Ritz
University of Wollongong, Australia

Ian S. Burnett
RMIT University, Melbourne, Australia

This article focuses on Internet video content and social networking to present solutions to the problem of gathering metadata that describes user interaction, usage, and opinions of that media.

There is an increasing quantity of video content on the Internet with popular services such as YouTube and Flickr allowing users to upload video material freely and easily. These services have resulted in user-generated content becoming a popular and everyday part of Internet culture. Simultaneously, the Internet has seen the rise of social-networking sites and activities with users increasingly incorporating media content on their personal pages. While most online content services started with closely guarded, proprietary interfaces, third-party access to content and data stored on these services via public APIs is becoming increasingly common. This article focuses on video content and social networking through these public APIs to present solutions to the problem of gathering metadata that describes user interaction, usage, and opinions of video content.

In this article, we focus on the description of video content and move beyond the simple tagging of a video as an object, which has proven successful and popular,[1-4] toward describing the actual temporal content inside the video. However, there still are lessons to be learned from the current object tagging used on most sites. On YouTube, video tags are currently limited to the author's selection, which presents issues because a video can generate different reactions from different users and user groups that might tag that video quite differently. Hence, general tag-based search of videos that evoke such disparate reactions will generally be unsatisfactory. In contrast, other problems arise when a video service allows all users to tag content; the result can then be many different tags with no indication as to which are relevant, reasonable, or correct. This can happen for the simplest of videos. For instance, one user might think a clip is funny whereas another finds it sad. Hence, while video sites have introduced the ability to search on the basis of simple tags, the semantics of the video for one user might not suggest that tag or even a closely related tag.

As video content has become increasingly popular, social-networking sites (such as MySpace, Facebook, and Multiply) have provided platforms that facilitate users building large groups of friends who can message, chat, and share media (from sites such as Flickr and YouTube). Just as with video sites, social-networking sites have exposed their information and interfaces as public APIs, but their aim has been to encourage the creation of novel applications. Facebook has led the way in this area, and as Facebook users create friend networks and interact with video and media content, they display implicit and explicit behavior that can lead to useful semantic data related to and describing that content. This information is exposed through the APIs. In this article, we discuss how we can collect this data for the semantics it reveals regarding the interaction between users, their friends, the groups to which they belong, and the media content. Our investigations show that when video content is being used and tagged in a social-networking context, there are many sources of semantic metadata available, much of it that would be beyond the reach of direct user tagging.

To explore the cauldron mix of media and social networks, this article explains the creation of the Tag!t application, which gathers semantics from social-group information and
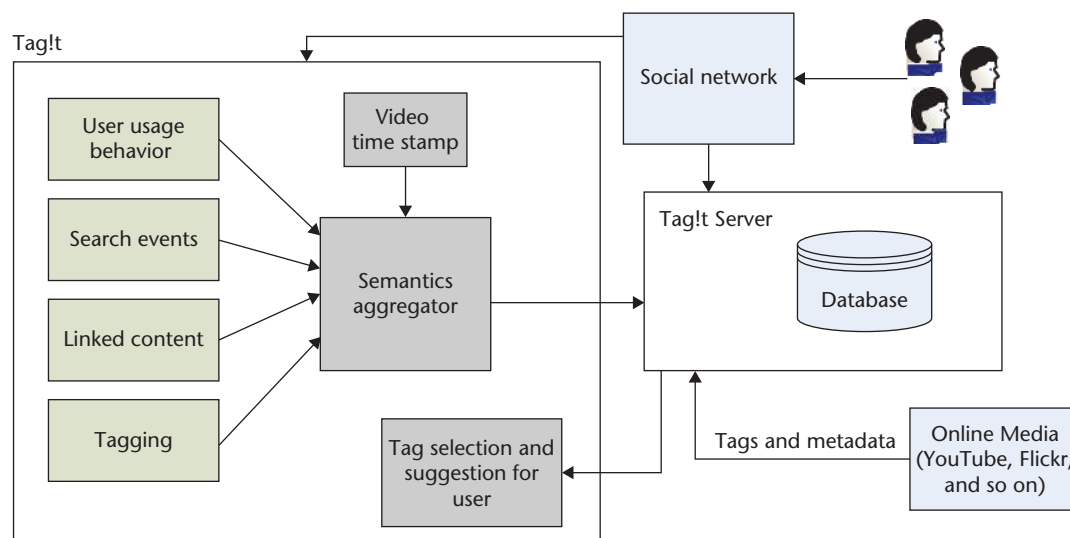
Tag!t

Internet media (illustrated in Figure 1). By recording the interactions of users with online media content, we acquire four types of semantics: user collection semantics, temporal tagging semantics, user behavioral semantics, and semantics from linked content. We aggregate these semantics and (for temporal media) the semantic event time stamps to form a collection of media content with semantic links constructed through social networking interactions. Tag!t could potentially improve search results for search engines (using social information) and make query languages (such as SPARQL[5]) more useful, and thereby improve the user's experience with the content.

Semantics and collaborative tagging

We can gather semantics from automated systems (for example, analysis tools) or manual approaches (for example, users tagging the content). Significant research already exists on extracting both low- and high-level semantics on the basis of video, audio, and subtitle analysis.[6-8] However, the majority of these techniques are targeted at specific search domains; a generic automated system would need to incorporate all these methods, and more, to create a rich and feature-complete set of semantics across many search domains. In contrast, manual tagging by users exploits background and collective knowledge built upon everyday experience. However, that richness of experience results in a great variety of tag sets[2] and the nature of those tags is highly dependent on users' incentives and motivations for tagging.[3-4]

Often it's difficult to persuade users to volunteer to add any semantics via tagging, and tagging users will do so for different reasons. In our work, we exploit social networks and groups to discover user and social group interactions. We can use this data to generate more consistent and less ambiguous tagging semantics for video content. By leveraging social networks, we can rely on the strength of many users contributing tags to a piece of content. We refer to this process as *collaborative tagging*. Social networks bring to collaborative tagging social relationships and interactivity, which lead to new tag characteristics and tagging behaviors.[9] Other research suggests that user incentives when tagging in social networks include future retrieval along with the ability to contribute and share, attract attention, play and compete, self-present, and express opinion.[3] Our work explores and exploits the results of these incentives for video and associated content.

In our collaborative-tagging system, users can tag either on an object-by-object basis or temporally. Typically, users within social groups collaborate because they can see the tags of others and add tags at the same time (to agree or disagree). Recently, open APIs have emerged to expose temporal tags so as to improve collaboration between services and social networks. Examples of these APIs are Annotea[10] and Co-Annotea,[11] which support associations between annotations and tags of mixed-media objects. Unfortunately, collaborative tagging currently suffers from two main problems: inconsistency and ambiguity. Both of these problems can be attributed to the polysemy, synonymy,

and haziness of user-generated tags.[2,3,9] However, researchers have observed that stable tags start to emerge when a substantial number of users tag content over a long period; we can think of this process as achieving a user consensus (effectively a standardization of terms by the group).[12] Other work has suggested that such consensus can be achieved by allowing users to correct other users' tags in way similar to the operation of wiki pages.[9] In this manner, social networks could build user consensus and more reliable semantic metadata for content.

Semantics from user interaction and usage behavior

In this section, we review the mechanisms by which we can derive semantic information about video content (illustrated in Figure 1) on the basis of combined user reactions and behaviors.

Usage scenario

To help illustrate the concept of semantic collection in this article, we present a simple usage scenario. Consider a user, let's call him Bob, who has just added a video from YouTube, called *The Original Human Tetris Performance*, created by Guillaume Reymond, into Tag!t. Because Bob has a Facebook account, the Facebook news feed alerts Bob's friends that he has added a new video. While watching the video, Bob is reminded of the old *Tetris* game he used to play on his Nintendo console. Bob enters the search terms ''*Tetris*'' and ''Nintendo'' into the search field and searches YouTube. Bob finds a relevant clip and decides to add it as part of the linked content. Bob repeats these steps searching Flickr for relevant pictures.

One of Bob's friends, Alice, is alerted by the Facebook news feed to the human *Tetris* video and decides to view it. Alice watches the video twice in its entirety while tagging selected scenes within the video. She is reminded of the real life Japanese human *Tetris* game from TV and adds links to some relevant videos. When Alice finishes her session, Tag!t uses the Facebook news feed to alert her friends that she has watched and tagged the human *Tetris* video. The story continues with the viral spread of the video, user tags, and linked content. But for this scenario, the scene is set.

Collection semantics

There are currently several approaches for describing collections or groupings of online multimedia content. One is the playlist approach, which is widely adopted in media players. We regard a playlist as simply a list of media that a user has decided to collect under a common theme (for example, a music or movie genre). Playlists described using XML (generally accepted as the language for cross-platform interoperability) include Advanced Stream Redirector;[13] Apple's iTunes library; and even standardized approaches using XML, such as XML Sharable Playlist Format.[14] In this article, we explore a broader set of media than that provided by the playlist model. Our approach is closer to the one used in MPEG-21,[15] which was created to describe the big picture of multimedia and multimedia interaction and provides a framework for grouping and anchoring relevant content and metadata together.

When users group content in a manner of their own choice, we are provided with extra semantics for that content. These semantics can be built from analysis of the collection name, other media added to that group, and the tags from the original media (for example, from YouTube). For instance, if Bob added the *Tetris* video to a collection called *stop motion*, we can deduce that the video is most likely a stop-motion video. If collection semantics are combined across many users, all placing content into containers, a result of, say, 90 percent of users within the same social group adding the same clip to a funny collection could lead to a high certainty that users in that group found the video genuinely funny.

In contrast, the same clip might end up in hundreds of people's collections, but with different chosen classifications; this might indicate different user opinions and connections generated by the video clip. In this vein, Figure 2 illustrates how a video about the funniest bloopers of all time can end up in different collections, but can be in the same collection within particular user social groups. Thus, the nature of social groups and the interactions within them can lead to improved collection tagging and semantic consistency for media collections.
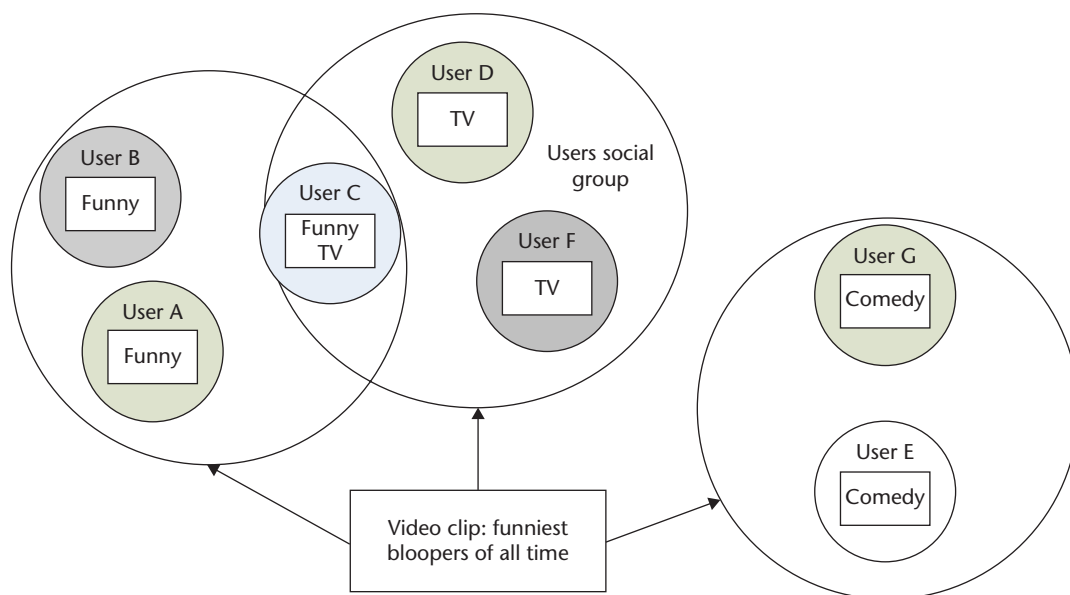
Temporal tagging semantics

The idea of temporal tagging is to allow users to add video tags that are anchored to particular time stamps. Popular sharing sites such as YouTube and Viddler have recently begun offering users the ability to tag temporally (or time annotate). Unlike tags that describe the entire clip, temporal tags might indicate the user's current feelings toward the current scene (for example, bored, laughing, and so on). Some users might wish to tag objects or people in that scene (for example, chair, John, and so on). And some users might wish to leave a scene comment.[16] Figure 3 illustrates the concept of temporal tagging and overlapping temporal tags.

In our usage scenario, Bob might have temporally tagged his clip with ''funny,'' ''awesome,'' and ''John Doe,'' and Alice might have added the ''awesome'' and ''boring'' tags. From these tags, relationships can emerge; for example, Figure 3 shows an overlap between ''boring'' and ''John Doe.'' If these relationships frequently happen, we can conclude that scenes with John Doe present are boring. Our database might show that this relationship is only present within certain social groups, and we could use additional information about these groups to discover more about the relationship. One application of the semantic information extracted from temporal tagging is the automatic creation of new mashup clips that combine segments of many different clips. For example, a user might like to see a video that contains all funny segments from the latest top 10 TV shows.

For temporal tagging to appeal to users and draw participation, it needs to be quick and simple (we demonstrated this in the initial testing of Tag!t, and other work highlights this observation[3]). Our experience has been that using one-click tags from a predetermined set is the best approach for the majority of users; anything more (for example, full comments) is only appropriate for particular presentation scenarios, such as for education. We have found that when it comes to user motivation for tagging, it's unusual for general users to feel the need to add substantial temporal comments to a video unless they are required to do so.

User behavioral semantics

We gather user behavioral semantics from the way users interact with the media, for example, monitoring the player's seek, pause, play, and stop events. While user behavioral monitoring has recently been added to YouTube (YouTube Insight), that service doesn't
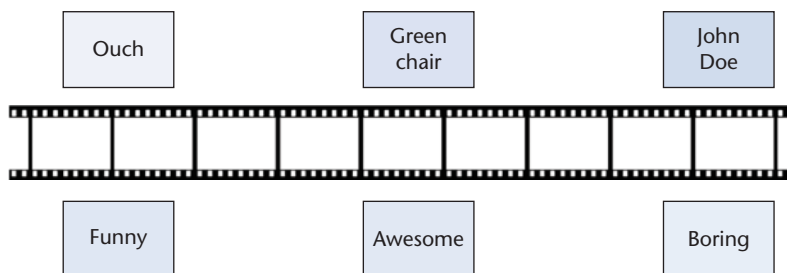
take advantage of the user's social groupings, as proposed here.

Comparing user behavioral semantics with other semantic metadata (for example, temporal tagging) can be used to generate a significant amount of extra information. For example, if Bob and some of his friends tagged a segment of clip as ''boring'' (temporal tagging), and numerous other people from the social group skipped the same segment, it's likely that the semantic of ''boring'' applies to that segment of media within that social group (and possibly other social groups too). This information could be used to automatically determine popular content (or segments of content), automatically remove video segments that cause users to constantly skip forward (thus ensuring users receive the most relevant part of the content), or improve search ranking results by ranking media clips that users don't watch in their entirety.

As previously discussed, these relationships might only be present within certain social groups. Thus, the search rankings could be targeted to users on the basis of their social groups. Doing so would be made feasible by the open API approach now offered by social-networking sites such as Facebook.

Linked content semantics

Referring to our usage scenario, we remember that while Bob was watching his video he was reminded of other content he thought was related to the video. By allowing users to add links to external content—that is, when users are reminded of related videos, photos, Web pages, and so on—we can extract semantic information about the relationship between content by examining the tags on the linked data. For example, when Alice links the Japanese real-life *Tetris* videos and Bob links the Nintendo *Tetris* videos, the following new tags supplied with the video provide extra semantics: Japanese *Tetris*, *Tetris*, human *Tetris*, computer, Game Boy, NES, and Nintendo. These tags are in addition to the author tags collected from the original YouTube video: stop motion, *Tetris*, play, real, people, human, live, seat, chair, sit, bleachers, art, interactive, performance, video, and game. This method of collecting semantics can provide us with a rich data set for further and deeper analysis.

While some users will directly link to content, others start a search for extra material as a result of viewing a particular temporal section or event in a video. By recording the search terms and the selections made at particular time stamps, we can build clues as to the meaning of that segment of the video clip. Additionally, we can determine how many people queried the same (or equivalent) term, so as to improve reliability and confidence in the tags. Such data lends itself to being a basis for suggesting temporally relevant material to new users.

Implementation

To validate the concepts of derived semantics presented in the previous section, we developed the Tag!t social-networking application. Tag!t is written in Flex and is embedded in a Facebook application that is designed to extract semantics both explicitly (that is, through direct user input) and implicitly (that is, automated without user input). The Tag!t system consists of two parts: a client-side application and a server, as illustrated in Figure 1.

The server—which we configured in a Linux, Apache, MySQL, and PHP (commonly known as LAMP) environment—receives the semantics from the Tag!t application and stores them in a MySQL database. When a user adds videos and additional content (typically from within Facebook pages), the server retrieves and stores the tags from those online media services for that content using the published APIs. For example, when a user adds a YouTube video, the server uses the YouTube data API to query for tags associated with that YouTube video.

As a user selects a video to view within Tag!t, the server dynamically constructs a Continuous Media Markup Language (CMML)[17] document containing the time points and tags generated by previous users for this video. Because CMML is dynamically created (that is, by selecting tags from the database), users can choose to restrict the received tag set to those tags selected from their friends and social groups. CMML gives us the ability to ask questions at given time points and even change the tag set throughout the video. This ability is mostly useful within educational environments and for conducting experiments. The CMML file is delivered to the client, which is then parsed and used to display the tags at the specified time within the video.

Tag!t builds upon existing multimedia sharing sites. Users are asked to place their videos

into a collection in order to group related videos. For example, a user might create a funniest clips collection. Such collections organize the videos to make it easier for the user and his or her friends to find content (and keep the interface similar to what users are accustomed to in Facebook, such as Facebook photos). Theses collections also are used by our system to extract collection semantics. After users add a video to their collection, they are taken to the main application, as illustrated in Figure 4. Tag!t facilitates a simple one-click button, temporal-tagging mechanism (providing temporal semantics) and connects related content to videos. As users watch the video, the tags are displayed at the nominated time stamp. All tags in the video can be viewed by clicking the expand button; double-clicking the tag moves the video to that particular time stamp.

Figure 1 illustrates the semantics that the Tag!t application extracts. Figure 4 shows the Tag!t application as seen on a user's Facebook page. Our emphasis here is not on the user interface, but rather the ways in which the application extracts the user's behavioral semantics. Tag!t observes and records the following:

■ *Tagging.* The tagging panel in Figure 4 lets users add temporal tags. Throughout the video, users can click an emotitag, create user-defined tags, and tag their Facebook friends. Emotitags are just emotion tags, while user-defined tags can be any tag of the user's choice. Under the video panel, the temporal tags are displayed at the corresponding time in the video.

■ *Search events.* The application provides a search box (see quick search in Figure 4) that can search YouTube and Flickr for additional content to connect to the current clip. Examples of additional content are relevant for similar videos, photos, and so on. The search terms and time in the video are stored because they might provide further insight about a particular scene.

■ *Linked content.* Users can link relevant photos and videos (resulting from search events) to the extras panel that will remain linked to that video (see the extras panel in Figure 4).

■ *User usage behavior.* The application monitors user behavior, such as pausing, seeking,

Figure 4. Screenshot of the Facebook application, Tag!t.
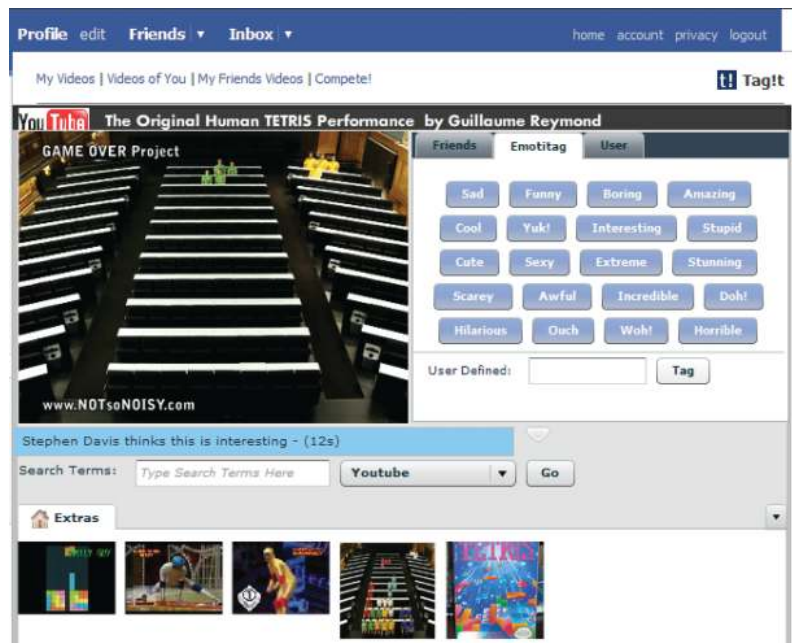
and exiting; the aggregation of these behaviors provides valuable information about the content within the video.

All of these behavioral semantics trigger a Tag!t semantic event that results in the semantics aggregator (see Figure 1) attaching the current video's time stamp to the semantic event and sending it to the server to be stored in the database. Recording the time of the semantic event facilitates analysis of the correlation between the video scenes and semantic events. On the basis of the set of semantic events for a video, the server can analyze the data to generate usage patterns, interest groups, and so on, and even provide suitable tags for suggestive tagging.[3]

Motivating users to tag can be challenging. Thus, to entice users to provide temporal tags, Tag!t offers a competition mode. Tag!t competitions allow users to compete with their friends by matching temporal tags and linking additional content. For the system demonstrated here, users aren't shown the current tags or linked content for a video and are asked to tag and link content as they see fit. Points are awarded on the basis of the temporal proximity to another friend's or user's tags. For example, if a user tags a video at precisely the same time as another user, then he or she is awarded more points than by adding the same tag some number of seconds apart. Points are also awarded for searching with the same search terms used by others.

Figure 5. Comparison of tag distribution over time between the technical and nontechnical groups.
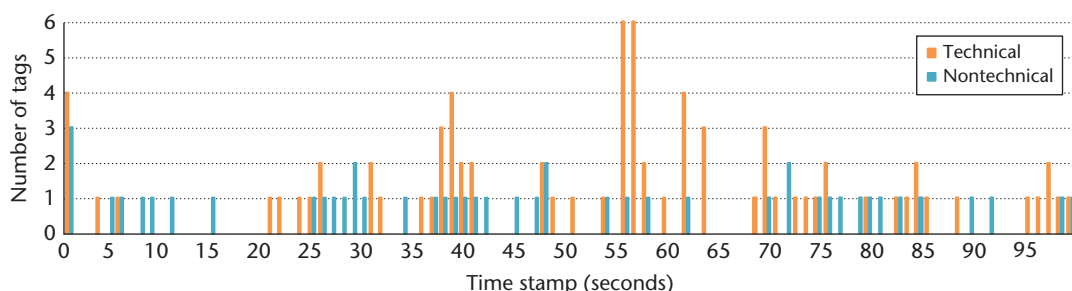
Evaluation

To evaluate semantics acquisition using the Tag!t system, we ran an experiment in which we asked users to tag eight video clips ranging from 68 seconds to 124 seconds. The video included genres such as comedy, documentary, news, sports, science, and music, with some videos selected to have distinct segments (such as a top five compilation) and other single segments. We had 22 users participate in the invitation-based experiment and asked them to select a primary group on the basis of their classifying themselves as technical or nontechnical—59 percent classified themselves as technical and 41 percent as nontechnical. In addition, we asked users to select several secondary groups in which they would be likely to participate on social-networking sites.

For the experiment, we limited the users choice of tags to 20 emotitags, which expressed emotions such as happy, sad, and boring. Users could add extras (other videos, pictures, and so on) at any time. We didn't allow users to create new tags because, as mentioned previously, we have found that creating new tags is mainly useful for educational video. For half the videos we used CMML to define time points within the video where the video would pause and prompt the user to tag or add an extra. It's important to note that users could still tag at any other point if they wished. For the other parts of the videos, we did not impose any forced tagging. Additionally, we configured four videos to hide other user's tags to see if that had any effect on when people would tag. When the users completed each video, we asked them to categorize the video.

Experimental observations

We collected a total of 860 tags and 133 extras across the eight videos in the experiment. Most notably, we found that there were distinct differences between the behavior of the technical and nontechnical groups. Figure 5 illustrates the tag distributions for technical and nontechnical users viewing one of the video clips (a short comedy scene involving physicists). In this case, users couldn't see other user tags and were free to tag at will. Figure 5 illustrates clustering of tags that directly correlate with the physics-related jokes within the clip, while Figure 6 shows a histogram of tags inserted by users for that clip. We saw these correlated group-clustering results across the majority of videos, including those that didn't prompt users to tag at specific times.

We also found that having other users' emotion tags hidden or visible during the clip did not have a significant impact on the number of tags contributed during a clip, which suggests that the quality of the content, rather than other user behavior, is the prime motivator for user tagging. This finding was also confirmed by our discovery that certain clips did entice users to search for relevant content and link those content objects as extras. Further, for the videos where we used CMML to prompt users to add extras, we found that many users within the same group entered the same (or semantically similar) search terms. We also found that some segments of video inspired users to search for music clips as opposed to video or picture content. The technical and nontechnical groups also showed differences between the search terms—that is, technical users searched for scientific terms, in contrast

Figure 6. Tag histogram for both technical and nontechnical groups for the video in Figure 5.
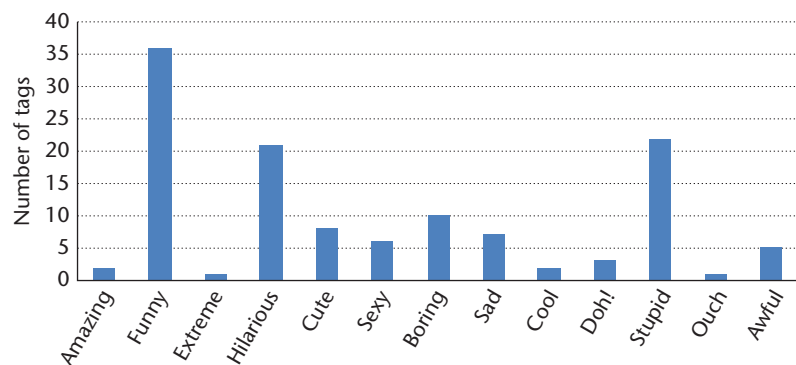
to more general terms searched by nontechnical users—when adding extras.

Overall, we found that the nontechnical users tended to add fewer tags; on average, there were 33 tags per nontechnical user as opposed to 42 tags for technical users. This disparity suggests that the interface is important if we are to persuade general, nontechnical users to become more interactive as they tag online videos.

Conclusion

In this article we have considered some of the many types of semantics that can be gathered from observation of user interaction with media content. In particular, we can exploit social networking to contextualize metadata added to and derived from media. Doing so is useful for improving the relevance of multimedia search results as well as for dynamically adapting the content presented to users according to usage patterns. Our future work will involve conducting larger-scale experiments to collect a large volume of semantics. Through the analysis of these semantics, we aim to investigate how tagging differs between social groups, consider mechanisms to automate suggested tags for users, and determine deeper semantics from combined user interactions. We expect these investigations to lead to new ontologies, based on user reaction to media, that will improve media-search algorithms and automatic repurposing of video. **MM**

References

1. M.G. Noll and C. Meinel, ''Authors vs. Readers: A Comparative Study of Document Metadata and Content in the WWW,'' *Proc. 2007 ACM Symp. Document Engineering,* ACM Press, 2007, pp. 177-186.
2. S.A. Golder and B.A. Huberman, ''Usage Patterns of Collaborative Tagging Systems,'' *J. Information Science,* vol. 32, no. 2, Apr. 2006, pp. 198-208.
3. C. Marlow et al., ''HT06, Tagging Paper, Taxonomy, Flickr, Academic Article, ToRead,'' *Proc. 7th Conf. Hypertext and Hypermedia,* ACM Press, 2006, pp. 31-40.
4. M. Ames and M. Naaman, ''Why We Tag: Motivations for Annotation in Mobile and Online Media,'' *Proc. Sigchi Conf. Human Factors in Computing Systems,* ACM Press, 2007, pp. 971-980.
5. E. Prud'hommeaux and A. Seaborne, eds., *SPARQL Query Language for RDF,* World Wide Web Consortium (W3C), 2008; http://www.w3.org/TR/rdf-sparql-query/.
6. H. Naphide and T. Huang, ''A Probabilistic Framework for Semantic Video Indexing, Filtering, and Retrieval,'' *IEEE Trans. Multimedia,* vol. 3, no. 1, Mar. 2001, pp. 141-151.
7. W. Zhou, A. Vellaikal, and C.C.J. Kuo, ''Rule-Based Video Classification System for Basketball Video Indexing,'' *Proc. ACM Workshop on Multimedia,* ACM Press, 2000, pp. 213-216.
8. C. Wu, Z. Chuang, and Y. Lin, ''Emotion Recognition from Text Using Semantic Labels and Separable Mixture Models,'' *ACM Trans. Asian Language Information Processing* (Talip), vol. 5, no. 2, 2006, pp. 165-183.
9. Q. Li and S.C. Lu, ''Collaborative Tagging Applications and Approaches,'' *IEEE MultiMedia,* vol. 15, no. 3, 2008, pp. 14-21.
10. J. Kahan and M. Koivunen, ''Annotea: An Open RDF Infrastructure for Shared Web Annotations,'' *Proc. 10th Int'l Conf. World Wide Web,* ACM Press, 2001, pp. 623-632.
11. J. Hunter and R. Schroeter, ''Co-Annotea: A System for Tagging Relationships Between Multiple Mixed-Media Objects,'' *IEEE MultiMedia,* vol. 15, no. 3, 2008, pp. 42-53.
12. A. Mikroyannidis, ''Toward a Social Semantic Web,'' *Computer,* vol. 40, no. 11, 2007, pp. 113-115.
13. Microsoft, *An Overview of Advanced Systems Format,* 2003; http://msdn.microsoft.com/en-us/library/ms983668.aspx.
14. L. Gonze, M. Friedrich, and R. Kaye, *XML Shareable Playlist Format Version 1,* 2008; http://www.xspf.org/xspf-v1.html.
15. I.S. Burnett et al., *The MPEG-21 Book,* John Wiley and Sons, 2006.
16. D. Yamamoto et al., ''Video Scene Annotation Based on Web Social Activities,'' *IEEE MultiMedia,* vol. 15, no. 3, 2008, pp. 22-32.
17. S. Pfeiffer, C. Parker, and C. Schremmer, ''Annodex: A Simple Architecture to Enable Hyperlinking, Search, and Retrieval of Time-Continuous Data on the Web,'' *Proc. 5th ACM SIGMM Int'l Workshop Multimedia Information Retrieval,* ACM Press, 2003, pp. 87-93.

Stephen J. Davis is a research fellow at the School of Electrical, Computer, and Telecommunications Engineering at the University of Wollongong, where he is funded by the Smart Services Technology CRC. His research interests include multimedia delivery, social networking and collaboration, multimedia semantics, and Web 2.0+. Davis has a PhD in computer engineering from the University of Wollongong. Contact him at stdavis@uow.edu.au.

Christian H. Ritz is a senior lecturer within the School of Electrical, Computer, and Telecommunications Engineering at the University of Wollongong. His research interests include single and multichannel speech signal processing; spatial audio coding; and multimedia content analysis, annotation, and delivery. Ritz has a PhD in electrical engineering from the University of Wollongong. Contact him at critz@uow.edu.au.

Ian S. Burnett is a professor and head of the School of Electrical and Computer Engineering at RMIT University, Melbourne, Australia. His research interests include multimedia processing and delivery, speech and audio coding, 3D spatial audio, and audio separation. Burnett has a PhD in electrical and electronic engineering from the University of Bath, UK. Contact him at ian.burnett@rmit.edu.au.

IEEE ⊕ computer society

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.

MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.

COMPUTER SOCIETY WEB SITE: www.computer.org

OMBUDSMAN: To check membership status or report a change of address, call the IEEE Member Services toll-free number, +1 800 678 4333 (US) or +1 732 981 0060 (international). Direct all other Computer Society-related questions—magazine delivery or unresolved complaints—to help@computer.org.

CHAPTERS: Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.

AVAILABLE INFORMATION: To obtain more information on any of the following, contact Customer Service at +1 714 821 8380 or +1 800 272 6657:

- Membership applications
- Publications catalog
- Draft standards and order forms
- Technical committee list
- Technical committee application
- Chapter start-up procedures
- Student scholarship information
- Volunteer leaders/staff directory
- IEEE senior member grade application (requires 10 years practice and significant performance in five of those 10)

PUBLICATIONS AND ACTIVITIES

Computer: The flagship publication of the IEEE Computer Society, *Computer*, publishes peer-reviewed technical content that covers all aspects of computer science, computer engineering, technology, and applications.

Periodicals: The society publishes 14 magazines, 12 transactions, and one letters. Refer to membership application or request information as noted above.

Conference Proceedings & Books: Conference Publishing Services publishes more than 175 titles every year. CS Press publishes books in partnership with John Wiley & Sons.

Standards Working Groups: More than 150 groups produce IEEE standards used throughout the world.

Technical Committees: TCs provide professional interaction in over 45 technical areas and directly influence computer engineering conferences and publications.

Conferences/Education: The society holds about 200 conferences each year and sponsors many educational activities, including computing science accreditation.

Certifications: The society offers two software developer credentials.

For more information, visit www.computer.org/certification.

◆**IEEE**

Celebrating 125 Years
of Engineering the Future

revised 1 May 2009

EXECUTIVE COMMITTEE

President: Susan K. (Kathy) Land, CSDP*
President-Elect: James D. Isaak*
Past President: Rangachar Kasturi*
Secretary: David A. Grier*
VP, Chapters Activities: Sattupathu V. Sankaran†
VP, Educational Activities: Alan Clements (2nd VP)*
VP, Professional Activities: James W. Moore†
VP, Publications: Sorel Reisman†
VP, Standards Activities: John Harauz†
VP, Technical & Conference Activities: John W. Walz (1st VP)*
Treasurer: Donald F. Shafer*
2008–2009 IEEE Division V Director: Deborah M. Cooper†
2009–2010 IEEE Division VIII Director: Stephen L. Diamond†
2009 IEEE Division V Director-Elect: Michael R. Williams†
Computer **Editor in Chief:** Carl K. Chang†

* voting member of the Board of Governors   † nonvoting member of the Board of Governors

BOARD OF GOVERNORS

Term Expiring 2009: Van L. Eden; Robert Dupuis; Frank E. Ferrante; Roger U. Fujii; Ann Q. Gates, CSDP; Juan E. Gilbert; Don F. Shafer
Term Expiring 2010: André Ivanov; Phillip A. Laplante; Itaru Mimura; Jon G. Rokne; Christina M. Schober; Ann E.K. Sobel; Jeffrey M. Voas
Term Expiring 2011: Elisa Bertino, George V. Cybenko, Ann DeMarle, David S. Ebert, David A. Grier, Hironori Kasahara, Steven L. Tanimoto

EXECUTIVE STAFF

Executive Director: Angela R. Burgess
Director, Business & Product Development: Ann Vu
Director, Finance & Accounting: John Miller
Director, Governance, & Associate Executive Director: Anne Marie Kelly
Director, Information Technology & Services: Carl Scott
Director, Membership Development: Violet S. Doan
Director, Products & Services: Evan Butterfield
Director, Sales & Marketing: Dick Price

COMPUTER SOCIETY OFFICES

Washington, D.C.: 2001 L St., Ste. 700, Washington, D.C. 20036
Phone: +1 202 371 0101 • **Fax:** +1 202 728 9614
Email: hq.ofc@computer.org
Los Alamitos: 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314
Phone: +1 714 821 8380
Email: help@computer.org
Membership & Publication Orders:
Phone: +1 800 272 6657 • **Fax:** +1 714 821 4641
Email: help@computer.org
Asia/Pacific: Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan
Phone: +81 3 3408 3118 • **Fax:** +81 3 3408 3553
Email: tokyo.ofc@computer.org

IEEE OFFICERS

President: John R. Vig
President-Elect: Pedro A. Ray
Past President: Lewis M. Terman
Secretary: Barry L. Shoop
Treasurer: Peter W. Staecker
VP, Educational Activities: Teofilo Ramos
VP, Publication Services & Products: Jon G. Rokne
VP, Membership & Geographic Activities: Joseph V. Lillie
President, Standards Association Board of Governors: W. Charlton Adams
VP, Technical Activities: Harold L. Flescher
IEEE Division V Director: Deborah M. Cooper
IEEE Division VIII Director: Stephen L. Diamond
President, IEEE-USA: Gordon W. Day

Next Board Meeting:
17 Nov. 2009, New Brunswick, NJ, USA