Social Tagging Revamped: Supporting the Users' Need of Self-promotion through Persuasive Techniques

Mauro Cherubini, Alejandro Gutierrez‡, Rodrigo de Oliveira, and Nuria Oliver
Telefonica Research
via Augusta, 177. Barcelona, Spain
{mauro, oliveira, nuriao}@tid.es
‡ University of Illinois at Urbana Champaign
201 North Goodwin Avenue Urbana IL 61801 USA
agutier3@illinois.edu

ABSTRACT
People share pictures online to increase their social presence. However, recent studies have shown that most of the content shared in social networks is not looked at by peers. Proper metadata can be generated and used to improve the retrieval of this content. In spite of this, we still lack solutions for collecting valid descriptors of content that can be used effectively in the context of social information navigation. In this paper, we propose a mechanism based on persuasive techniques to support peers in providing metadata for multimedia content that can be used for a person's self-promotion. Through an iterative design and experimentation process, we demonstrate how this methodology can be used effectively to increase one's social presence thus building more enjoyable, rich, and creative content that is shared in the social network. In addition, we highlight implications that inform the design of social games with a purpose.

Keywords: Information overload, Facebook fatigue, metadata, mutual modeling, self-presentation, social networks

ACM Classification Keywords: H.5.3 **[Group and Organization Interfaces]**: Collaborative computing

General Terms: Human Factors

INTRODUCTION
People use social networking sites for two main reasons: to monitor what contacts and friends are doing and as a self-presentation tool [11]. Two features of social networking sites are particularly relevant to our research: (1) they allow their users to keep a constant level of awareness on their peers' activities; and (2) they support a one-to-many communication style which targets trusted members of the social network. A typical way of participating in this communication flow involves sharing multimedia content with peers. All the major social networks now offer the ability to their users to share pictures and videos with their friends and family members. This feature has become more popular by the day such that users transfer increasing amounts of their user generated content from their personal repositories to their social networks. For instance, Facebook photos is the largest consumer contributed photo-sharing service in the world in terms of users and is considered one of the fastest-growing of any size with more than 2 billion photos uploaded to the site each month [1].

As a result of the continuous growth of shared online content, most of the users of social networks are overloaded by constantly changing feeds and they struggle to keep up with the content. Although the phenomenon has not been addressed extensively in the scientific community, it is a well known issue in the online community and has been referred to as "Facebook fatigue" [5]. As a result, multimedia content spreads slowly throughout the social network and it is often overlooked by peers [7]. In short, information overload is hindering the social awareness and self-presentation needs that users try to satisfy in social networks.

A number of interaction mechanisms (see background section for more details) are available in current social networking sites that help users cope with this information abundance. Nevertheless, both the automatic summarization of the content (*e.g.*, the Wall in Facebook) or manual mechanisms such as *social filtering* [14] and *social navigation* [16] seem to be inefficient solutions to the problem because: (1) They do not lower the effort required by the user to fulfill the needs described above because their implementation still requires the user to go through their peers' feeds–at least in existing implementations in social networking sites–; and (2) they do not prevent users from loosing sight of valuable content that might be available in their social network at a given time.

In this paper, we propose the use of *persuasive* techniques [9] to elicit knowledge (in the form of metadata or tags) from peers in the social network that can be used to support and promote the user's self-presentation. In addition, the output of these techniques might allow new interaction strategies with multimedia content that could reduce the burden of keeping up-to-date with the peers' activity.

We look at *social tagging* [22] from a new perspective: its current definition focuses on the collaborative aspect of the creation and management of metadata to organize and annotate user generated content, independently of the social relationship between the participants and the owner of the content. In our work, we focus on –and structure the activity within– networks of peers that know each other (*i.e.*, social networks). As friends and family members possess mental models [21] of each other, they are in a privileged position

CHI 2010, April 1015, 2010, Atlanta, Georgia, USA.

to create rich descriptors of the user-generated multimedia content.

In the next section, we summarize the background research that inspired this work and pose the research questions that we address in this paper. Next, we describe the iterative methodology that we applied to identify the factors that play a role when eliciting metadata from the social network. We then present the user studies that we carried out to validate our findings. Finally, we discuss the results of the study and propose several implications for the design of social tagging applications.

BACKGROUND RESEARCH

Social Networks, Social Capital, and Appearance

Social networks are used to maintain relationships with friends and to meet new people. Joinson investigated the uses of the social networking site Facebook (FB) and the types of gratification that users derive from those uses [11], probing more in depth the exact nature of "keeping in touch", as both use and gratification. The results of his work suggested that the "keeping in touch" use comprises two main functions: a *surveillance* and a *self-presentation* function. Concerning the latter, he found that many users did not change their privacy settings from the default configuration. Making one's profile more accessible increases the chances of an interaction. Associated with this use is the gratification derived from building "social capital", where FB is used to build and maintain ties with peers [8]. Social capital here is defined as *the sum of the resources, actual or virtual, that accrue to an individual or a group by virtue of the social network* [6]. The feedback provided by social partners onto one's online networking profile affect the observer's impressions of the profile's owner [25]. These findings are relevant to address the question of why people share digital content through social networks.

Why Do People Share Pictures Online?

Miller and Edwards interviewed a group of Flickr[1] users and found that power users considered tagging to be a social activity where they could include inside references and jokes as tags [17]. Conversely, normal users tagged infrequently because they could easily retrieve their pictures using their chronological order. About 50% of the participants interviewed reported sharing pictures through web sites such as Flickr or via email. Furthermore, Van House *et al*. found that sharing pictures with remote peers is a way of maintaining relationships, while telling stories with these photos helps nurture the relationships [23].

Digital photography has significantly lowered the costs of taking and sharing pictures to the point that users can easily flood their acquaintances with content that is less "photo-worthy" than it used to be. Miller and Edwards confirmed previous findings that shared pictures were used to tell stories *with* rather than *about* them [17, p. 2]. An interesting finding of the study was that current photoware tools did not support the storytelling aspect of photo sharing that was so important to the "Kodak culture". Different stand-alone solutions have been proposed in the past for supporting storytelling with digital pictures (one of the seminal papers in this area is the study of Balabanović *et al*. [3]). However, additional research needs to be carried out to determine whether these findings still apply in the context of social networks, and to study how to support sharing practices within social networks.

[1]See `http://www.flickr.com`, last retrieved September 2009.

Information Overload and Strategies to Cope with Complexity

Social networks are largely a non-monetary environment. However, they follow economic rules. They are at the root of a "crowding out" effect as lower quality content crowds out good quality content by way of very "low prices" (*i.e*., large availability) for online content. Content shared through social networks may give raise to adverse selection [2]. In other words, without a proper mechanism to distinguish good quality from low quality content, owners of good content may be reluctant to publish their content thus lowering the average quality of all shared content. A solution to adverse selection is *signaling* [2]: Producers should be able to signal the quality of their work through an objective and unambiguous signal. However, we still lack effective solutions for the producer of content to better advertise the quality of his/her contributions. Similarly, we lack effective solutions for the consumer of this content to readily inspect the quality of the material s/he can get access to.

While adverse selection has still to be demonstrated to take place on social networks, researchers are already observing slow information propagation in social networks. For instance, Meeyoung and colleagues collected and analyzed traces of information dissemination in a large social network [7]. They found that even popular pictures do not spread widely and rapidly throughout the networks. They also found that information is usually exchanged between friends with a significant delay at each hop.

In addition to these attempts to objectively measure information overload, little is known on the techniques that users put in place to cope with this abundance of information. We can speculate that they might restrict their focus to a specific group of trusted friends that pre-process the information for them (*i.e*., social filtering [14]) by reading and reposting relevant news. Alternatively, we might expect users to take advantage of social navigation traces that are offered by most social networking sites (*e.g*., the preference ranking, sometimes referred as "thumbs up/down"). Moreover, most of these systems offer comprehensive tagging functionalities that could lower the barriers of retrieving multimedia content. However, we know little about whether other mechanisms exist and are used. More importantly, these mechanisms are generally suboptimal solutions to the problem because they require manual effort from the users in order to navigate, organize, and aggregate information. A possible solution would be to enrich multimedia content with more descriptive metadata through advanced tagging techniques, as discussed next.

Social Tagging and GWAP

The study of Miller and Edwards [17] revealed that people tag infrequently or not at all. Also, people tend to organize their pictures by date and events. Hence, their retrieval, in most situations, did not seem problematic. Users did realize that a potential benefit of tagging is that it can aid in finding pictures more easily. However, users add tags to multimedia content not only to facilitate their retrieval but also to describe images to family and friends [18]. Also, researchers have found evidence that tagging might be used to increase social presence [18, 22]. There are, however, a few shortcomings intrinsic to user-generated tags: (1) a lack of standardization of the chosen vocabulary; (2) the effort

that is necessary to collect them; and (3) a lack of precision in choosing the best words that could describe the content. Concerning this third point, a recent study conducted by Marques and Lux [15] revealed that 46.2% of tags used to describe pictures in Flickr are not useful to properly categorize and describe the contents of the images.

These limitations of user-generated tags have motivated scholars to find alternative methods to generate descriptors of multimedia content. An interesting approach in this domain is the *ESP Game* [24], an online game where players that do not know each other see the same picture and type words that describe the content of the picture. When the same word is typed by the two players, the word is considered to be *valid* and it is used as a tag for the picture. The game has been shown to be effective in generating a large amount of valid tags in a short amount of time. Von Ahn later designed similar games that he named Games With A Purpose (GWAP). Recently, a study by Robertson *et al*. [20] measured a shortcoming of the ESP game: namely the *redundancy* in the tag sets. That is, often synonyms are provided as tags of the same picture (*e.g*., "man" and "guy"). Of all 496 (out of 14.5K) images labeled as "guy", 81% were also labeled as "man"[2]. In other words, the very game mechanic encourages players to enter generic labels (*e.g*., "building" as opposed to "terraced house") in order to maximize the chances that the same word will be entered by the other player.

Robertson *et al*.'s study made us wonder whether designing a game with people that know each other could yield more specific and useful descriptions of the multimedia content.

Social Tagging Revamped: Persuasive Techniques and Mutual Modeling

Social tagging is currently defined as the activity of producing collaboratively metadata in the form of keywords to shared content [10]. In this paper, we propose a new definition of the term social tagging where the "social" term refers to social networks. We believe that this shift of focus might have two benefits:

1. *The tagging activity might be supported by the mutual model of the peers*. Mutual acquaintances possess accurate information about each other (*e.g*., the work they do, where they live, what they like, and so forth). This information is organized into a mental model that is usually referred to as *mutual* because of its reciprocal nature [21][3]. Our hypothesis is that this mutual model might help produce more descriptive metadata for the multimedia content because of the implicit knowledge that the peers in a social network might have about the content of a friend.
2. *The tagging activity might be corroborated by persuasive techniques*. Persuasive techniques are interactive computing products created for the purpose of changing people's attitudes or behaviors [9]. Most of these techniques are based on social influence theory, like social compliance or social support. In particular, the natural context in which these persuasive techniques are mostly effective is within a social network because of the psychology of human relations [12]. Our hypothesis is that by contextualizing the tagging activity to the social network, persuasive techniques can be designed that would encourage people to provide richer metadata for the multimedia content.

There are a few examples of this social dimension of tagging in the literature. Quian and Feijs designed a system to annotate pictures while chatting [19]. Their system could extract information from conversations to generate metadata for the shared pictures. However, their system was challenged by limitations of natural language processing and the abbreviations and jargon usually used in chat conversations. Similarly, Barthelmess and colleagues designed a semi-automated labeling application based on the extraction of metadata from naturally occurring conversations of co-located people looking at pictures [4].

Based on our definition of social tagging, the research presented in this paper addresses the three research questions stated below.

Summary and Research Questions

One of the main reasons why people use social networks is to self-promote themselves. Often, this is achieved by sharing multimedia content. This leads to an overload of information published by peers in SNs. To date few tools are available to the users to contrast this information abundance. The study of Joinson [11] did not provide evidences of the fatigue that users are subjected to. Therefore, we attempt to provide more systematic evidences of this phenomena by posing the following question: *RQ1 – How do people cope with information overload in social networks (Facebook)?*.

One of the common solutions to this problem is to have valid descriptors for the content (*i.e*., metadata). These can be obtained using automatic or crowd sourcing techniques. However these methods have drawbacks. We propose to exploit the implicit knowledge that peers, from the same social network, have of each other (*i.e*., mutual model) to produce this metadata. To verify this could be a viable solution, we posed RQ2: *When describing the same multimedia content, do members of the same social network generate terms that are more specific than those generated by people outside the social network?*

Finally, Akerlof suggested that a signaling functionality can be useful to people to distinguish information sources [2]. Therefore, we thought that the descriptors obtained within the social network could be used to design a signaling functionality for multimedia content shared among the peers (*i.e*., promote a peers profile in a social network). We also wondered whether such signaling functionality could be readily appropriated by the users. To verify these ideas we posed RQ3: *Can we design a quality control mechanism (involving persuasive techniques) such that social network peers would be willing to provide meaningful descriptors of multimedia content that would support personal promotion?*. Next we present an overview of our experimental plan.

METHODOLOGY

We report the results of three studies that were designed to shed light on each of the previously presented research questions. To answer RQ1, we designed and deployed a questionnaire to three social networks in Facebook (with an average size of 16 peers). To answer RQ2, we enrolled

[2]These statistics were computed using a freely available dataset of images tagged with ESP. See, `http://bit.ly/7tcYyR`, last retrieved January 2010.

[3]*Mutual modeling* is the process of building assumptions about the beliefs and the knowledge of one's partner(s). Sangin *et al*. [21] demonstrated that this process is grounded at group level: *i.e*., A posses a model of B and B possess a model of A. These might not be symmetrical.

9 employees of a large corporation. They were invited to share pictures and interact with each other's photos while we collected their qualitative feedback. Finally, we designed a semi-controlled experiment to measure whether commentaries generated through structured interactions with peers – such as the ones that took place in our second study – could be used for personal promotion in Facebook (FB). The results of this last study were used to validate RQ3.

STUDY 1: EXPLORATORY QUESTIONNAIRE

Participants

The participants on the first study were 48 Facebook users (m:36, f:12 with a median age of = 26.5 years) who responded to an email request to complete a short online study which was open during the second week of September 2009.

They belonged to three independent social networks (SN): (a) the first SN was of graduate students, U.S. residents of different nationalities, different ages, and areas of study; (b) the second SN had both students and full-time employees of a large software company specialized in customer services. Subjects were on non-immigrant visas while in the U.S. and were from different nationalities; (c) participants of the third SN consisted of an international group of researchers and administrative assistant within a big telecommunications company in Spain.

Materials

The online survey comprised basic demographic questions (*e.g.*, age, gender, etc), alongside some measures of their Facebook use (*e.g.*, time spent on FB, number of friends on FB, history of use, etc.). Next, participants were asked to respond to the following questions using free text:

A. What do you normally use Facebook for?
B. Why do you share pictures on social networking sites, if any at all?
C. What are the strategies that you use to keep up with what your friends are doing on Facebook?

We also included 5-point rating scale items and multiple choice questions in the questionnaire for the following items:

D. Do you feel overwhelmed by the amount of updates of your peers' activity in Facebook?
E. How often do you look and click on the posted pictures of your Facebook friends?
F. Which of the following factors play an important role to you when determining whether you will look at a picture or not?

Procedure

We collected data from the three different social networks, where their members (17, 15, and 16 participants respectively) did not know any of the members in the other social networks. For each of the open-ended questions, we defined a coding scheme based on the most recurring typologies of answers. Two researchers independently coded the entries for each question and there was a nominal disagreement (1%). These differences were resolved via discussion.

Results

The first goal of the analysis of the questionnaire was to understand how respondents used social networking sites, whether they shared pictures or not, and what strategies people had in place to stay informed of their friends' activity.

Participants had an average of 277.9 friends linked to their Facebook profile (Range $10 - 1600$, Median $= 200$, $s = 307.8$). Two of the participants had been registered on the site for less than six months. The remaining participants had been signed up for between six months and a year (6.2%), more than one year, but less than two (29.2%) or for more than two years (58.3%). The majority of participants visited the site either daily (34.1%) or more than once a day (25.0%). Fourteen visited FB several times a week (31.8%) and only 4 participants (corresponding to 9.1%) visited their FB page less than once a month.

When asked what they normally use FB for (item A above), the most common answer (12 respondents, or 25%) was for sharing pictures and other multimedia content with friends. The second most common use (11 respondents, or 22.9%) was for maintaining relationships with distant friends, followed by as a communication tool (10, or 20.8%) and to maintain perpetual contact with friends (9, or 18.8%). Finally, only 4 respondents (or 8.3%) declared using FB to reacquire lost contacts, and 2 (or 4.2%) as a way of watch what others were doing. These results were consistent with the findings of Joinson [11].

Then we asked whether and *why* they shared pictures on FB (item B above). Forty respondents (83.3%) reported sharing pictures regularly, while 4 (8.3%) respondents declared sharing only occasionally, and 4 (8.3%) did not share at all. From those who share pictures, 21 respondents (52.5%) declared sharing to stay close to friends and family members, 11 (27.5%) to show the visited places, holidays, and events, and 4 participants (10%) reported using FB as an easier distribution method than email. The last 4 participants declared using it because of the availability of the tagging and commenting features. These findings were consistent with the study of van House and colleagues [23].

Coming to the question that addresses **RQ1** (item C above), respondents reported using three basic solutions when dealing with the constantly changing information on sites like FB: (1) checking the feeds often (23 respondents or 48%), (2) using manual filters (11 respondents or 23%), and (3) using automatic filters (4 respondents or 8%).[4] Ten respondents (21%) reported not having any strategy to deal with information overload.

We looked at item D to check whether information overload was perceived as a problem by the respondents. Four participants (8.3%) reported feeling *always* overwhelmed and another 4 felt overwhelmed *often*. The majority of respondents (24 or 50%) declared having the problem *sometimes*. Eight people (16.7%) declared perceiving the load *rarely* and another 8 declared *never* perceiving the problem. These results confirm that information overload is a problem in social networking sites. Most of the respondents felt overwhelmed by the amount of updates in FB and they adopted few strategies to overcome it. When asked whether they felt there was a lack of tools to help them deal with their friends' updates, 39.6% of the participants that perceived the load (68.4%) answered *Yes*.

[4]By manual filters we mean that participants declared to parse visually their feeds looking for specific categories, like picture posts. Conversely, by automatic filters we mean that participants used advanced features of their social networking site to restrict the updates to a subgroup of their peers.

Additionally, item E in the questionnaire asked respondents how often they interacted with their friends' multimedia content. Seven participants (14.6%) reported looking at their friends' pictures every day to several times per day. Twenty participants (41.7%), declared looking at the pictures several times per week and 13 (27.1%), looked at them a few times per month. The last 8 (16.7%) respondents declared looking at pictures only once per month.

Finally, when asked what factors played a role in determining what content to interact with (item F), respondents selected: *the owner of the picture* (40 times); *the appeal of the picture* (34 times); *the number of comments* (10 times); *the content of the comments* (10 times), and *the number of likes and dislikes* (9 times).

STUDY 2: PLAYING WITH MUTUAL MODELING

Participants

Nine subjects (m: 8, f: 1) were recruited by email advertisement in a large corporation. Their ages ranged from 23 to 46 years ($\bar{x} = 31$, $s = 7.76$) and they were Computer Science researchers, graduate students and administrative assistants working part or full-time. The study (comprising three phases) was open from the beginning throughout mid-June.

Procedure

The study was divided into three phases. In the first phase, participants were asked to submit a photo from their personal repository that they thought would be of interest to the other participants taking place in the experiment (colleagues). On the following day, participants were asked to describe – via a Web interface – their colleagues' pictures with nouns, adjectives and verbs. These tags could contain multiple (*e.g.,* ping pong, group meeting) or single (*e.g.,* ball, people) words.

In the second phase, subjects were requested to type a comment about their friends' pictures using a similar web interface to the first phase. Finally, in the third phase, each participant saw the tags and comments that his/her colleagues had typed for his/her picture and rated each one as *"I liked it"*, *"I didn't like it"*, or *"Neutral"*. Furthermore, they provided explanations for why they liked or disliked the tags and comments, and also submitted final thoughts about the whole experiment.

Results

This qualitative user study confirmed general concepts regarding social networks that are usually considered as common sense. We highlight next the most relevant findings in order to address **RQ2**.

People tend to like photo comments from their peers, mostly when they include jokes. After taking part in this experiment, participants made it clear that reading the peers' comments was a pleasurable experience. *E.g.*, *"I REALLY liked to see my friends' comments on my picture"* (user 2), *"User 9*[5] *comments were pretty good"* (user 8). Conversely, analyzing the answers provided by participants to the questions of whether they found it useful to look at the tags that their friends provided for their picture, we could not find evidence of agreement nor disagreement. Table 1 shows the number of comments that each photo owner liked or did not like to read and the main reasons why.

[5]The original sentence included the name of the participant, which was anonymized.

Table 1. Number of comments that photo owners' liked and did not like to read and main reasons why

Photo	# comments that the photo owner		Main reason for the photo owner	
	liked	*didn't like*	*liking a comment*	*not liking a comment*
1	1	2	Know peers better Owner: *"Interesting to see how some of my colleagues still do not know where I fly from after a year"*	Low Quality Owner: *"Not useful"*
2	2	1	Joke Owner: *"I liked the joke"*	Low Quality Owner: *"No creativity"*
3	5	0	Joke Owner: *"I like this comment because it's funny :)"*	—
4	5	0	—	—
5	2	3	Joke Owner: *"he he...true, but this is from France... I guess the culture carries over"*	Tease the peer Owner: *"Who the hell has this for breakfast? What's wrong with you!"*
6	2	1	Joke Owner: *"funny comment"*	Low Quality Owner: *"too many questions and sentences are too long"*
7	6	1	Joke Owner: *"good joke!"*	—
8	2	2	—	Tease the peer Owner: *"'that guy' is ambiguous. 'acting' implies motion and is misleading in this context."*
9	4	2	Good Quality Owner: *"curiosity ... good for later retrieval"*	Low Quality Owner: *"I know that"*

According to Table 1, jokes seem to be the main reason why participants liked to read their peers' comments. Furthermore, we noticed that a two of the subjects declared to not like a few comments because they wanted to write a teasing reply to the peer (see rows 5 and 8 in Table 1). This behavior could also be considered as a kind of joke. These observations support the following conclusion.

Commenting is typically a communication activity directed to a person, while tagging is impersonal. Many sentences posted by the subjects at the end of the experiment directly contribute to this statement: *"I think that keyword based tagging is better for content retrieval because the addressee of the communication is the anonymous world and thus the terms are often chosen in order to explain the picture. As for the comment, I directed my communication to the author of the image"* (user 4), *"To some of these comments I would have liked to write a little follow-up"* (user 1).

Relationships between peers are different and affect comment appreciation. The different levels of friendship among peers was made clear even with the small social network considered in this study (nine peers), as can be noticed from the following comments: *"Knowing who the owner of the picture is may result different comments or tagging"* (user 6), *"I know the people I work directly with very well and understand their personality and sense of humor. I don't know newer people or interns as well and maybe I'm more inclined to like comments/descriptions by people I know and understand. Also, perhaps I could misinterpret something written by someone I know less well"* (user 3).

RQ2 – When describing the same multimedia content, peers belonging to the same social network generate terms that are more specific than those generated by people outside the social network. Taking into consideration the

previous observations, we addressed **RQ2** by manually classifying all the tags generated by the subjects into two groups: (1) *general tags*, or tags that are obvious by observing the picture itself; and (2) *specific tags* or tags that include knowledge of the owner of the picture and the context of the social network. For example, a photo with five men sitting around a table had the tag "group meeting" classified as a general tag and "Monday group meeting" as a specific tag, as the picture does not provide temporal information, whilst the members of the social network know that the picture was taken during a group meeting that usually occurs every Monday. Table shows the number of general and specific tags used by people inside and outside the social network in order to retrieve each of the nine photos.

Table 2. Comparison of general and specific tags used by people inside and outside the social network to retrieve the photos used in Study 3.

Photo	Inside SN (peers)		Outside SN (strangers)	
	# unique general tags	# unique specific tags	# unique general tags	# unique specific tags
1	24	5	24	4
2	28	5	26	1
3	24	11	22	2
4	24	2	21	0
5	24	1	22	0
6	22	0	23	0
7	22	2	21	1
8	17	7	19	3
9	22	4	14	2

According to Table , the number of general tags is larger than specific tags for both people inside and outside the social network. We believe that this is due to the fact that all subjects –and the majority of today's computer users– are familiar with web search engines (*e.g.*, Google, Yahoo, *etc.*). Hence, they do not expect that using specific tags related to the author of the picture or the corresponding social network could enhance their chances of retrieving it at a later time. Table also shows that peers inside the social netork generate more specific tags than strangers do, thus supporting the validity of **RQ2**. Perhaps more important is the observation that most of the specific single-word keywords were contained in the commentaries that participants wrote during the second phase. For example, one of the participant, named Raul (fiction name) posted a picture of himself with a glass of wine. Peers used the word "Raul" as single-word tag and one of the participants also used the word "Raul" in his commentary: (participant 2) *"Raul is drinking cheap wine!!!"*. Therefore, this finding suggested the idea that commentaries could be used as metadata instead of single-word tags.

Tagging and commenting has potential beyond search and retrieval. As previously presented, the comments by peers were mostly appreciated when they improved the user experience in social networks by making it more enjoyable and fun. In addition, peers generate more specific tags than people outside the social network. Thus, this finding suggests that there are opportunities for reconsidering the process of tagging within the social network realm. We describe next the user study that we conducted to further explore this venue of research.

STUDY 3: SELF-PROMOTION EXPERIMENT

Participants

Fifty-one Facebook users (m:40, f:11) were recruited by e-mail advertisement and belonged to three different social networks (SN): (1) SN1, with 17 participants (m:15, f:2) that were mostly graduate students from several areas of study, had different nationalities, were all U.S. residents, and their ages ranged between 23 and 30 years ($\bar{x} = 27, s = 2.06$).

(2) SN2, with 14 participants (m:11, f:3) that were both students and full-time employees of a large software company specialized in customer services. Subjects were non-immigrants in the U.S. and were from different nationalities. Their ages varied from 22 to 47 years ($\bar{x} = 28, s = 8.22$).

(3) SN3, with 20 participants (m:14, f:6) that were mostly Computer Science researchers within a big telecommunications company in Spain, had different nationalities, and ages varied from 23 to 46 years ($\bar{x} = 31, s = 5.99$). The study was open from the final week of August until the second week of September.

Procedure

In Study 3 we designed two quality control mechanisms for the multimedia content to be shared through the social network. The first method (*PhotoBest*) consists of three steps: 1) the user uploads an album; 2) some peers vote for the best pictures of that album; and 3) the three most voted pictures of the album are published on the main feed to advertise the content of the album and its contributor. The second method, (*CommBest*) consists of four steps: 1) the user uploads an album and chooses one representative picture for that album; 2) some peers are asked by the contributor to create a funny commentary for the leading picture; 3) the contributor selects the best commentary; and 4) the chosen commentary and associated picture are published on the main feed to advertise the content of the album, its contributor, and the author of the winning comment. Picture 1, represents these two mechanisms visually.

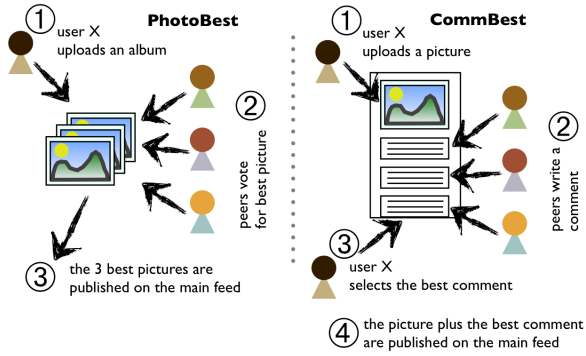

Figure 1. The two quality signalling mechanisms tested in study 3

Four subjects (photo providers)[6] from each social network were asked to send us 50 photos from one of their online photo albums and to choose the one photo that they thought could better promote their album. This material was used to initiate the two mechanisms described above.

As soon as this preparation phase was over, the actual self-promotion experiment was conducted as follows. First, the remaining peers of each social network (SN1: 17, SN2: 14,

[6]These four subjects from each SN were used only to gather data for the actual self-promotion experiment and they were not part of the figures presented in the Participants section. If we were to include them, the total size of the sample would be 63 instead of 51.

SN3: 20 participants) were asked to access four web pages similar to the one presented in Figure 2 and to choose in each page the feed that they would choose in order to further explore the photo album. Note that this interface imitates the Facebook style for presenting feeds, which was the social network commonly used by all participants in the experiment.Each of the four web pages presented four different feeds promoting the *same* photo album from one of the four *photo providers* from the preparation phase.

The four different feeds correspond to the multimedia presentation techniques that were evaluated in this study: (1) *PhotoRnd* – three photos chosen randomly from the album, (2) *CommRnd* – one photo and associated comment selected randomly, (3) *PhotoBest* – best three photos as selected by the *PhotoBest* quality control mechanism, and (4) *CommBest* – the photo and associated comment as selected by the *CommBest* quality control mechanism presented above.

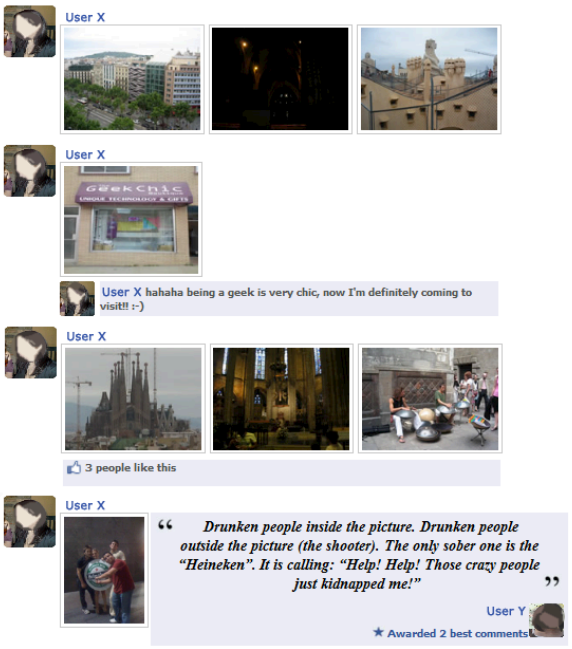

Figure 2. Web interface of the self-promotion experiment. Four different multimedia presentation techniques in Facebook style were evaluated: Random photos (*PhotoRnd*), Random comment (*CommRnd*), Best photos (*PhotoBest*), Best comment (*CommBest*). The faces and the names of the participants have been anonymized

The presentation order of the four different techniques followed a Latin square model[7] to avoid biasing the results. Users' interaction was logged and once they confirmed their choice for a particular feed, they also provided us with their main reasons for selecting that feed presentation mode for each webpage.

Results

The three social networks considered in this experiment differ in their size (SN1: 17; SN2: 14; SN3: 20 participants), subjects' country of residence (SN1 and SN2: United States; SN3: Spain), nationalities (*i.e.*, mixed), and domains of expertise (*e.g.*, Computer Science, Psychology, Medicine, *etc.*).

[7]See `http://tinyurl.com/latinsq`, last retrieved September 2009.

Although such diversity in the samples increases the study's external validity, it is important to check whether it had a significant effect over the observed variables among different social networks. We carry out such an analysis next.

Subjects from different SN's took the same time to finish the trials: No significant difference was found in the time spent by the subjects from different social networks to read each of their friends' feeds and choose the one they would like to further explore ($N = 204, \chi^2 = 1.527, df = 2, p = .47$). This analysis was done without grouping the duration of the four trials per subject, and therefore it reveals the real task of browsing one's profile. Furthermore, we found no correlation between the duration of the trials and the subjects' preferred method ($N = 204, \rho = -.084, p = .24$).

Subjects from different SN's clicked on feeds the same number of times: ($N = 204, \chi^2 = 2.516, df = 2, p = .28$). Additionally, there was no correlation between number of clicks in the trials and the subjects' preferred method ($N = 204, \rho = .086, p = .17$). Table summarizes the main descriptive statistics for variables *trial duration* and *clicks per trial*.

Table 3. Study 3 descriptive statistics (Q_n: quartile).

SN	Trial duration (s)			Clicks per trial			Trials	Subj.
	Q_1	Q_2 ($\tilde{x}$)	Q_3	Q_1	Q_2 ($\tilde{x}$)	Q_3		
1	58	82	116	1	1	2	4	17
2	60	91	167	1	1	1	4	14
3	53	83	129	1	1	1	4	20

Subjects from different SN's had the same preferences for the presentation methods of the feeds: The methods chosen in all trials were the same among the three social networks ($N = 204, \chi^2 = 1.476, df = 6, p = .96$), and no significant difference was found on the preferred methods per trial regardless of the social network ($N = 204, C = .275, \chi^2 = 16.755, df = 9, p = .05$). The latter result also confirms that the presentation of the stimuli did not influence the subjects' choice. Table shows a crosstabulation of the preferred method and the social network, while Table crosses the preferred method and the trials.

Table 4. Crosstabulation of the subjects' preferred method and social network ($\chi^2 = 1.476, p = .96$).

Method	SN1	SN2	SN3	Total
PhotoRnd	16	16	22	54
CommRnd	8	4	8	20
PhotoBest	27	24	30	81
CommBest	17	12	20	49
Total	68	56	80	204

Table 5. Crosstabulation of the subjects' preferred method and trial number ($\chi^2 = 16.755, p = .05$).

Method	Trial 1	Trial 2	Trial 3	Trial 4	Total
PhotoRnd	15	14	12	13	54
CommRnd	2	8	2	10	20
PhotoBest	23	14	27	17	81
CommBest	11	17	10	11	49
Total	51	51	51	51	204

From these observations, we corroborate that the different social networks and the four trials considered in our study didn't lead to different results. Therefore, we combined these

statistics in order to draw some conclusions about the preferred multimedia presentation for self-promotion in social networks.

Introducing a quality control mechanism to select the *photos* that appear on a certain feed increases their persuasiveness: In order to reach this conclusion, we processed the data by computing the number of times each method was chosen by each user in his/her four trials (*e.g.*, if user 1 chose *PhotoRnd* in trial 1, *CommRnd* in trials 2 and 3, and *CommBest* in trial 4, the user's post-processed data would be *PhotoRnd*: 1, *CommRnd*: 2, *PhotoBest*: 0, and *CommBest*: 1). Figure 3 shows a bar graph with the users' preferred methods in percentage.

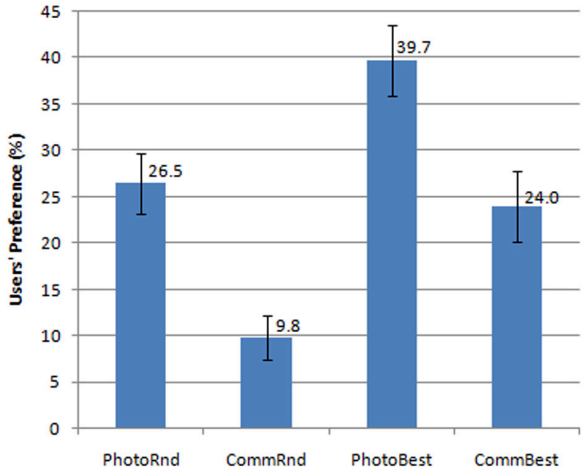

Figure 3. Users' preference for the best persuasive presentation method of multimedia content in social networks.

As can be noticed from the graph, there's a significant difference between the users' preferred methods ($N = 51, \chi^2 = 29.529, df = 3, p < .01$), particularly between the one with random photos –*PhotoRnd*– and its equivalent with the best photos as voted by other peers from the same SN –*PhotoBest*– ($N = 51, Z = -2.232, p = .03$). These results confirm that participants were more persuaded to explore feeds advertised with photos chosen by their friends (quality filter) than with random photos (39.7% against 26.5% of the user's preference).

Introducing a quality control mechanism to choose the *comment* that shall be presented on a certain feed increases its persuasiveness: Likewise, the error bar graph shown in Figure 3 heuristically confirm a significant difference between the users' preference for methods *CommRnd* and *CommBest* ($N = 51, Z = -2.519, p = .01$). This finding supports the previous one: comments that have been generated by means of an entertaining structure seem to attract more attention than those usually posted in social networks. After parsing the user's responses to the post-experiment questionnaire, it was clear that people really liked the idea of having quotes and jokes to represent an album (*e.g.*, user 3 from SN1: *"A cool comment can really make all the difference when I'm deciding what to look at on someone's feed."*, user 6 from SN1: *"Funny comment!"*, user 2 from SN2: *"I like how the person comment is put in quotations"*, user 17 from SN3: *"The star and the award of best comment caught my attention."*).

Introducing a quality control mechanism to choose the *comment* that shall be presented on a certain feed can be as persuasive as choosing random photos (no quality control): As can be noticed from the bars in Figure 3, photos are more persuasive than comments to promote one's photo album. However, the impact of introducing a quality filter in the comment's choosing process can highly improve it's appeal towards acquiring the user's preference. This is demonstrated by the fact that no significant difference could be identified between the user's preference for methods *PhotoRnd* and *CommBest* ($N = 51, Z = -.483, p = .63$).

Therefore, we answer **RQ3** by stating that applying a quality filter for both the photo selection and the comment generation processes can have a significant positive effect in the promotion of one's shared photo album.

GENERAL DISCUSSION

RQ1: How do people cope with information overload in FB?
The first study focused on the different strategies that people adopt to deal with information overload in social networking sites. Respondents reported using three solutions when dealing with the constantly changing information on sites like Facebook: (1) checking the feeds often (23 respondents, or 48%), (2) using manual filters (11 respondents or 23%), (3) and using automatic filters (4 respondents or 8%). Ten respondents (or 21%) declared not having any strategy to deal with information overload.

Note that all the strategies reported above are sub-optimal solutions to the problem of information overload because they do not effectively address the issue of missing relevant information that might appear in the feeds at a given time. Checking the feeds often or using manual filters requires an additional effort from the user who might feel overwhelmed (see results from user study 1). In fact, 68.4% of our participants declared perceiving the load of the constantly changing information. Concerning the automatic filters, users can effectively reduce the burden of going through lots of feed updates. However, this solution does not prevent users from missing relevant news that might be filtered out by the criteria imposed by the system.

These findings resonate well with the idea introduced by Akerlof [2] that the abundance of less relevant content might "crowd out" more relevant content. Given the rate at which social networks are growing and the rate at which content is increasingly shared through them [1], we believe the HCI community should pay more attention to designing more effective solutions that would allow users to signal the quality of the content uploaded in the network. For this reason, we conducted the second study reported in this paper.

Additionally, our sample reported reasons for using social networking sites that are consistent with the findings of Joinson [11]. Respondents declared sharing pictures in social networks for the same reasons they share them through email or face-to-face, as found by Miller and Edwards [17].

RQ2: When describing the same multimedia content, peers from the same social network generate terms that are more specific than those generated by people outside the social nework.
The second study demonstrated that peers in a social network provide a larger number of specific tags than those provided by strangers when describing a picture taken by one

of the members of their social network. This result is synergic with the idea expressed by Robertson and colleagues [20] that the interaction mechanics following which people generate metadata for a given content influences the quality of the descriptors that are produced. This finding is relevant for the design of tagging systems because it shows that users naturally take advantage of their implicit knowledge about their peers when tagging content provided by their peers. Additionally, the second study confirmed the findings of Miller and Edwards [17], and Nov *et al*. [18] who reported that users are not prone to tag their content because they consider it to be a tedious and boring activity whose benefits are not fully understood [13]. Similarly, our participants reported the activity of generating single word tags as being boring and uninteresting. Conversely, when we asked them to create comments for the pictures, they felt much more engaged and entertained. They liked the ability of creating conversations, (*i.e.*, stories), around pictures and being able to vote for their friends' best comment. Finally, we observed that by imposing a structure around the commenting activity, we obtained comments that are more relevant descriptions of the content when compared to the comments that are typically found in social networking sites. These results suggest that structured approaches to entering comments could leverage the fun component and the persuasive elements derived from carrying out the activity with friends. This qualitative observation was then confirmed in the third study.

RQ3: Can we design a quality control mechanism –involving persuasive techniques– such that social network peers would be willing to provide meaningful descriptors of multimedia content that would support personal promotion?
The main finding of the third study is that pictures that have been selected by peers to represent an album are better signals to promote one's multimedia content than random pictures taken from the album. Similarly, comments (of pictures) generated by the SN peers and through a structured activity are a better signal to promote multimedia content than standard commentaries. Participants explained that they liked the display of the picture plus best-rated commentary (*CommBest*) because: a) the comment made them think the album was going to be entertaining (*e.g.*, part. 3, sn 3, "*Looks like a fun night out! I want to find out more about it*."); b) they knew the author of the commentary and therefore they wanted to know more about the album (*e.g.*, part. 2, sn 2, "*A friend I knew made the comment*."); c) they liked the commentary (*e.g.*, part. 7, sn 1, "*The comments are my favorite part about people's pictures so I like to see that others commented on the pictures and it will make me want to look at the pictures*."). Additionally, this study demonstrated that when the comments are contextualized within the social network, they are able to attract as much attention as the standard random photos presented by most social networks to promote one's multimedia content. Therefore, our findings support the idea that the two quality control mechanisms that we designed are able to effectively signal the quality of the multimedia content and to support personal promotion.

Although it is easier to inspect the content of an album by looking at a sample of thumbnails, this study showed that contextualized comments might allow the users of social networking sites to gain access to interpretative information they could not gain elsewhere. Furthermore, the study demonstrated that the most attractive methodology is the one that combines the visual support of inspecting many thumbnails at once with social navigation (*PhotoBest*). Participants reported relying their decision on the fact that other peers had already inspected the quality of the multimedia content: (*e.g.*, part. 4, sn 1, "There are multiple people that I know in this content and I see that other people liked it for some reason.").

The major advantage of the approach we presented in this paper is that no advanced techniques are required in order to generate interesting descriptors for the multimedia content. It follows the same principle of GWAP [24]. This result is relevant for the design of tagging systems because it provides empirical support to the idea that by framing people's interaction in the social network we can stimulate people in creating richer metadata for their multimedia content.

IMPLICATIONS FOR DESIGN

There is a Need for Signaling Strategies

The first implication that this paper raises is that users of social networking sites need solutions to signal the quality of the multimedia content they deal with. This is true for both the producers (*i.e.*, those who upload the content) and the consumers of multimedia content. Effective signaling should allow consumers to sample the quality of the content without the need of inspecting the content. For instance, the number of "thumbs up" that appear in the bottom of a post allow efficient social navigation to which most of the users are accustomed to, as revealed by study 3 reported in this paper. Social navigation and social filtering are two of the mechanisms by which we can support information navigation in social networking sites. The commenting mechanism exploited in this study was based mostly on social support (*i.e.*, mutual help in a social network) because contributors asked their peers to create commentaries and their peers did it because of friendship. However, other powerful persuasive mechanisms are available to the wise designer: *social influence* [12], namely, when an individual's thoughts or actions are affected by other people, *entertainment* and exploitation of *personal benefits*.

Social Games With A Purpose

This work also suggests the raise of a new class of games with a purpose (GWAP) that would include two elements: *social networks* and *entertainment*. In the second study, we found that mutual modeling allows peers to contribute to tagging with knowledge that is not available to strangers. Therefore, this study suggests the possibility of creating games with immediate benefits for groups of people that know each other. We name this class of games: Social Games With a Purpose (SGWAP). We believe that the area of SGWAPs deserves more research as such games could enable novel forms of *entertaining* interactions between the users have in social networking sites while achieving useful goals, such as tagging multimedia content.

Tags Are Not Just for Retrieval

In this paper, we have used the terms *tag* and *comment* as synonyms for two reasons: First, we found in study 2 that users have an hard time describing content in terms of single-word statements; second, study 3 showed that comments can be used for the promotion of personal content and not only for retrieval purposes. We subscribe fully to the principle expressed by Kustaniwitz and Scheiderman [13] that in order to obtain metadata from the users, designers have to lower the barriers needed to input this information and to raise the incentives that people might get from contributing this information to the system. This study shows that by asking users

to input natural language commentaries instead of atomic tags, we can increase the reward that users get from the activity because they can tell stories through the commentaries [23], and interact and communicate with their friends. Furthermore, we believe that these commentaries might be also used to support information retrieval in combination with natural language processing engines. Of course, solutions in this field is still far from perfection, but they are worth investigating (*e.g.*, [4, 19]).

Take into Consideration Social Proximity

Social networks, such as those established in applications as Facebook, tend to be very large with users having hundreds of "friends" that hardly know each other and that lack the background knowledge to leave meaningful comments / tags (while the groups studied in the experiments reported in this paper were relatively small). Thus, a future improvement of the social tagging mechanism needs to consider the *social proximity* among peers, *e.g.*, by applying weights.

CONCLUSIONS

The main contributions of this work can be summarized in three points: 1) We have demonstrated how information overload is a real problem for which users do not have many efficient strategies at hand; 2) We have shown how by leveraging the knowledge of peers in social networks, it might be possible to obtain relevant descriptors for multimedia content that are complex and difficult to obtain with other techniques; and finally, 3) we have found that these descriptors might be used to signal the quality of multimedia content, thus supporting self-promotion.

We consider this work as a first step towards the definition of Social Games with a Purpose, a new class of GWAP that could take advantage of the specific properties of social networks and that could offer immediate benefits for people that know each other. Future work will require testing the applicability of these techniques to larger populations and to implement a game framework around the proposed techniques. Additionally, we plan to design other interaction mechanisms that could exploit other unique features of social networks as *social compliance*.

In conclusion, we believe that by turning the attention to the unique properties of social networks, social tagging might reveal possibilities unexplored before.

REFERENCES

1. Official statistics of facebook, `http://bit.ly/12oAN`, last retrieved January 2010.
2. AKERLOF, G. A. The market for lemons: Quality uncertainty and the market mechanism. *The Quartely Journal of Economics 84*, 3 (August 1970), 488–500.
3. BALABANOVIĆ, M., CHU, L. L., AND WOLFF, G. J. Storytelling with digital photographs. In *Proc. of CHI '00* (New York, NY, USA, 2000), ACM, pp. 564–571.
4. BARTHELMESS, P., KAISER, E., HUANG, X., MCGEE, D., AND COHEN, P. Collaborative multimodal photo annotation over digital paper. In *Proc. of ICMI '06* (New York, NY, USA, 2006), ACM, pp. 4–11.
5. BENENSON, F. Information overload, facebook fatigue, and twitter's awesome filter, `http://bit.ly/wlkA`, last retrieved January 2010.
6. BOURDIEU, P., AND WACQUANT, L. *An invitation to reflexive sociology*. University of Chicago Press, Chicago, USA, 1992.
7. CHA, M., MISLOVE, A., AND GUMMADI, K. P. A measurement-driven analysis of information propagation in the Flickr social network. In *Proc. of WWW '09* (New York, NY, USA, 2009), ACM, pp. 721–730.
8. ELLISON, N., STEINFIELD, C., AND LAMPE, C. Spatially bounded online social networks and social capital: The role of Facebook. In *Annual Conference of the International Communication Association (ICA)* (Dresden, Germany, June 19-23 2006).
9. FOGG, B. Persuasive computers: perspectives and research directions. In *Proc. of CHI '98*, (New York, NY, USA), pp. 225–232, ACM Press/Addison-Wesley Publishing Co., 1998.
10. GOLDER, S. A., AND HUBERMAN, B. A. Usage patterns of collaborative tagging systems. *Journal of Information Science 32*, 2 (April 2006), 198–208.
11. JOINSON, A. N. Looking at, looking up or keeping up with people? Motives and use of Facebook. In *Proc. of CHI '08* (New York, NY, USA, 2008), ACM, pp. 1027–1036.
12. KELMAN, H. C. Compliance, identification, and internalization: Three processes of attitude change. *Journal of Conflict Resolution 2*, 1 (January 1958), 51–60.
13. KUSTANIWITZ, J., AND SHNEIDERMAN, B. Motivating annotation for personal digital photo libraries: Lowering barriers while raising incentives. Technical Report HCIL-2004-18, University of Mariland, College Park, MD, USA, 2005.
14. LERMAN, K. Social networks and social information filtering on digg. In *Proc. of ICWSM'07* (Boulder, Colorado, USA, March 26-28 2007).
15. MARQUES, O., AND LUX, M. An exploratory study on joint analysis of visual classification in narrow domains and the discriminative power of tags. In *Proc. of MS '08* (New York, NY, USA, 2008), ACM, pp. 40–47.
16. MILLEN, D. R., AND FEINBERG, J. Using social tagging to improve social navigation. In *Workshop on the Social Navigation and Community-Based Adaptation Technologies (AH'06)* (Dublin, Ireland, June 20th 2006).
17. MILLER, A. D., AND EDWARDS, W. K. Give and take: a study of consumer photo-sharing culture and practice. In *Proc. of CHI '07* (New York, NY, USA, 2007), ACM, pp. 347–356.
18. NOV, O., NAAMAN, M., AND YE, C. What drives content tagging: the case of photos on Flickr. In *Proc. of CHI '08* (New York, NY, USA, 2008), ACM, pp. 1097–1100.
19. QIAN, Y., AND FEIJS, L. M. G. Exploring the potentials of combining photo annotating tasks with instant messaging fun. In *Proc. of MUM '04* (New York, NY, USA, 2004), ACM, pp. 11–17.
20. ROBERTSON, S., VOJNOVIC, M., AND WEBER, I. Rethinking the ESP game. In *Proc. of CHI'09* (New York, NY, USA, 2009), ACM, pp. 3937–3942.
21. SANGIN, M., NOVA, N., MOLINARI, G., AND DILLENBOURG, P. Partner modeling is mutual. In *Proc. of CSCL'07* (New York, NY, USA, July 16-21 2007).
22. THOM-SANTELLI, J., MULLER, M. J., AND MILLEN, D. R. Social tagging roles: publishers, evangelists, leaders. In *Proc. of CHI '08* (New York, NY, USA, 2008), ACM, pp. 1041–1044.
23. VAN HOUSE, N., DAVIS, M., AMES, M., FINN, M., AND VISWANATHAN, V. The uses of personal networked digital imaging: an empirical study of cameraphone photos and sharing. In *Proc. of CHI '05* (New York, NY, USA, 2005), ACM, pp. 1853–1856.
24. VON AHN, L., AND DABBISH, L. Labeling images with a computer game. In *Proc. of CHI '04* (New York, NY, USA, 2004), ACM, pp. 319–326.
25. WALTHER, J. B., HEIDE, B. V. D., KIM, S.-Y., WESTERMAN, D., AND TONG, S. T. The role of friends' appearance and behavior on evaluations of individuals on Facebook: Are we known by the company we keep? *Human Comm. Research 34*, 1 (January 2008), 28–49.