Recognition Markets and Visual Privacy

Ryan Shaw
UC Berkeley School of Information
ryanshaw@sims.berkeley.edu

Every year on April 20th, University of Colorado students gather in Farrand Field for an act of civil disobedience and hedonism: the mass consumption of marijuana. In 2006, the University of Colorado Police Department prepared for this event by placing cameras around the field and posting signs notifying visitors that all activity there would be photographed and videotaped. Confident in their numbers, the students ignored the signs, perhaps secure in their belief that though they might be seen, they wouldn't be recognized. The UCPD duly captured images of the ensuing conflagration, and then made a novel move: they posted the images on a public web server, with the explanation that anyone who successfully identified a pot-smoking miscreant would receive fifty dollars (Frauenfelder). Within days a substantial portion of those depicted had been recognized.[1]

The UCPD used the web to enroll a mass audience in the task of identification, in effect creating a human-machine hybrid technology of image recognition. Throwing images against a wall of eyeballs to see what sticks is not entirely new: consider the Most Wanted posters on the wall at the post office. What is new is the scale and speed at which the Internet can bring together those willing to sell their powers of recognition with those willing to buy. This realization has resulted in an explosion of proposals for, and some implementations of, new socio-technical systems which use "human-in-the-loop" approaches to solve image recognition problems. Often these systems are operationalized using economic models and treated as markets in which recognition services are exchanged for some form of capital. I will thus refer to such systems as "recognition markets." This paper will present an overview of the state of the art in recognition markets as well as possible developments in the future, discuss their implications for visual privacy, and suggest some guidelines for their ethical design and implementation.

Why recognition?

Why focus on recognition rather than more generally on surveillance? Image capture technologies—tiny cameras, zoom lenses, and cheap sensors—have become incredibly sophisticated and ubiquitous. But despite the rapid development and proliferation of image capture devices, there still seems to be a sense that we can "hide in plain sight" of the unblinking eyes that surround us. Perhaps we are comforted by the sheer volume of images being captured, secure in the feeling that we are needles in haystacks. In other words, maybe we don't mind being seen, as long as we aren't recognized. If so, this suggests that the discussion about visual privacy should focus not on image capture technologies, but on image recognition technologies. Recognition links images to specific people, places, and things. It is these links that allow us to use images to evoke memories, provide evidence, establish proof, or spread propaganda.

Signal processing for recognition

Traditionally, recognition systems have depended on signal processing. To electrical engineers, a

---

1   As witnessed by the author on the CU-Boulder Police Department's "420 Photo Album,"
    <http://www.colorado.edu/police/420_Photo_Album/>, last accessed April 29, 2006.

"signal" is any measurable quantity that changes over time or space. Audio is a one-dimensional signal: frequency varying over time. An image, on the other hand, is a two-dimensional signal varying in the horizontal and vertical directions. Video is a three-dimensional signal, varying spatially within any particular frame like an image, but also varying over time like an audio signal. The field of signal processing is based upon the fundamental notion that all of these types of signals can be represented as combinations of a basic set of mathematical functions. Thus signal processing provides a universal language for describing and manipulating what appear to us as very different things. All of the many different ways in which perceptual phenomena manifest themselves are, through the lens of signal processing, transformed into "features" which can be measured and compared. The goal of recognition is to build models which use the presence or absence of these perceptual features to determine whether an image depicts some entity or concept. These entities or concepts, things like "Vegetation" or "Corporate Leader," are often referred to as "high-level features," and the recognition task is thus called "high-level feature extraction."[2]

Signal processing-based recognition systems follow a canonical design pattern, which I present a very simplified description of here. First, images of the entity to be recognized are collected and analyzed as signals. Each individual image or sequence of images is then represented as a collection of the signal features to which it has been reduced. Next, these sets of features are used to build a statistical model of the entity. There are a variety of ways to build this model, but the end goal is a function that will, when presented with a new image or video as input, tell whether or not it depicts that entity (perhaps along with a number representing the level of confidence it has that its answer is correct). When applied to an archive of images, such a function can be used to select a subset of images depicting the desired entity.

A common metric for evaluating the accuracy of a recognition system is "precision at 100."[3] The system is asked for 100 images of some entity it has been trained to recognize, such as "Homeless Person/Hobo."[4] Then those images are examined by human evaluators, and the number of images that actually do depict the object is counted. This number is the system's precision at 100. If the system is perfectly accurate, the precision at 100 will be 100. If the system is just randomly guessing, the precision at 100 will depend on how many images in the archive depict the query object, since even a random guesser will be right occasionally, given enough chances. A recognition system is considered successful if it consistently performs better than random guessing.

All else being equal, the more input data used to build a recognition model, the more accurate it will be. Building an accurate model requires not just a large number of images, but a large number of *correctly labeled* images—in other words, images that have been verified to depict the object to be recognized. But not all objects are equally recognizable. It turns out that recognizing the face or body of a specific individual is very difficult, even with huge amounts of input data. Even for very commonly photographed individuals like Madeline Albright, the very best recognition systems have a precision at 100 of around 30 (Naphade). This is far better than chance, but still very far from human-level recognition accuracy. Better results can be achieved by constraining the problem in various ways, for example by restricting the set of images to full-face, front-view portraits captured indoors under controlled lighting conditions. But in terms of recognizing people in large archives of arbitrary images, signal processing-based systems still

---

2   See the "TRECVID 2006 Guidelines" <http://www-nlpir.nist.gov/projects/tv2006/tv2006.html> for a specific example of a high-level feature extraction test and how it is evaluated.

3   The evaluation metric presented here is somewhat simplified. For a more detailed explanation of the kind of metric actually used, see "Inferred Average Precision and TRECVID 2006" <http://www-nlpir.nist.gov/projects/tv2006/infAP.html>.

4   This example was taken from the LSCOM Lexicon <http://www.ee.columbia.edu/dvmm/lscom/>, a list of concepts developed for the evaluation of recognition systems by a government-industry research consortium.

have a long way to go (Lew et al).

Using volunteer labor for recognition

Purely signal-based approaches are not the only way to link images to the things that they depict. Humans are far better than machines at recognizing faces (though still far from perfect).[5] According to conventional wisdom, however, people are expensive. Much research into signal processing for recognition has been based upon the assumption that the recognition process needs to be fully automated to be feasible. People may be able to identify every picture of Bill Clinton without fail, but how long will it take them to examine thousands or millions of images? Better to have them examine relatively few— enough to train a statistical model—and leave the rest to the algorithms.

Such reasoning made sense before the spread of global, high-bandwidth communication networks. Now, however, it is possible to divide such tasks among millions of people. Harnessing large-scale cooperation among unpaid volunteers to construct information products in this manner has received a tremendous amount of attention in recent years. Law professor Yochai Benkler has coined the term "commons-based peer production" to refer to the phenomenon of flat, distributed networks of collaborators producing goods which they hold in common (Benkler). He argues that this form of production can outperform both markets and organizational hierarchies when the following conditions hold: first, that the "goods" being produced are information or culture, and second, that the means of producing these goods—computers and telecommunications networks—are widely distributed.[6] Benkler further claims that successful peer production systems have three characteristics. First, they must be modular, meaning the work to be done can split into pieces and handled incrementally and asynchronously. Second, the modules of work must be highly granular, meaning that participants can make even very small contributions to the overall effort. Finally, successful peer production systems must have automated or semi-automated systems for integrating contributions and exercising quality control.

The linkage of images to labels by masses of people can be seen as exhibiting all of Benkler's characteristics. Labels can be assigned to a given image in stages, by different people working independently. Photo-sharing and communication tools like Flickr[7] and Facebook[8] have made the linking and labeling of images and people and things quite easy, enabling even casual users to make contributions. Meanwhile there is considerable effort being put into the design of both systems of incentives for encouraging the contribution of "quality" information, and tools for identifying and weeding out "poor-quality" information. A few commercial media management systems such as Flickr appear to have successfully leveraged peer production of image labels to enhance the precision of their image retrieval. By allowing users to create labels that are useful for them personally, and then aggregating large numbers of these labels and making them (semi-) publicly available, these systems succeed both in encouraging a small minority of users to create labels, and in making those labels useful to many others.

In 1999 David Stork was among the first people to recognize the potential of combining peer production

---

5   For example, consider Levin's claim that people have difficulty recognizing faces with ethnic features to which they have not been regularly exposed.
6   I present Benkler's arguments here not to express my support for them, but to explain how the image recognition has been framed as a peer production problem. For a critical review of Benkler and others' claims for the superiority of peer production, see Duguid.
7   Yahoo! Inc., "Flickr – Photo Sharing," <http://flickr.com/>, last accessed October 29, 2006.
8   Facebook, <http://www.facebook.com/>, last accessed October 29, 2006.

with artificial intellingence or "machine learning," describing a system for harvesting and aggregating small bits of useful information from masses of people (Hearst et al). Luis Von Ahn later applied this concept to image labeling with the ESP Game (von Ahn & Dabbish). In this game, two anonymous players collaborate to describe an image. Both players are presented with the same image, and are asked to type labels that describe the image. Each player sees only her own labels. As soon as both players have entered the same label, they are considered to have "agreed" on that label, and the process begins again with a new image. The game is to find strategies that result in rapid agreement. A later variation, Peekaboom, uses a similar approach to label not just whole images but specific regions of images, making it possible for recognizers using those labels to recognize individuals in group photos (von Ahn et al). Recently Google has incorporated von Ahn's software into its image search engine.[9]

Human labeling of images can be combined with signal processing to accelerate the recognition process. Human-labeled images can be provided as training data for recognition models, and the results of human labeling can be compared to algorithmic results to check quality and make improvements (Muller et al). This technology is already making its way into the marketplace. Riya[10] and Polar Rose[11] are two recently formed companies hoping to use distributed annotations from their users to train recognition models, which can then be used to make the annotation process easier. The hope is that such systems will form a virtuous cycle, with the end result being a web of images as easily searchable as the text web is today.

The labeling of images with keywords representing concepts, as well as the algorithmic manipulation of images and their associated labels to enable search, are both forms of what Julian Warner calls "semiotic labor:" work done with and on symbols and signs (Warner). Warner categorizes semiotic labor into syntactic and semantic labor, where the former is concerned with formal representations of meaning, while the latter is concerned with meaning itself. He argues that the history of information systems has been to transform semantic labor into syntactic labor, suitable for algorithmic execution by computers. In the case of recognition systems, we see the opposite trend as well: syntactic systems replacing purely syntactic labor with hybrid systems that outsource their semantic labor to millions of humans. The digital economy has long relied on free labor (Terranova), so it is no surprise to see that systems for utilizing this labor are becoming thoroughly rationalized. Exchanges in which masses of workers who complete some micro-task are rewarded with money or entertainment, such as Amazon.com Inc.'s "Mechanical Turk"[12] and Mycroft Inc.'s "Mycroft Network,"[13] exemplify this new breed of labor market. Given that volunteer labor pools promise low-cost but bandwidth-intensive solutions to difficult problems, there is naturally a high level of interest in such systems. There has been a small but steady stream of academic papers devoted to the understanding of these systems [Kelly et al., Ludford et al., Beenen et al.], alongside a smaller, less steady trickle of critique.[14] Recently MIT announced the creation of cross-disciplinary center devoted to research into "collectively intelligent" assemblages of people and machines.[15]

Getting more personal

These systems can efficiently recognize common objects and well-known people and places, but should they be cause for concern among ordinary people? Consider an image of yourself, perhaps a candid photo

9   Google Inc., "Google Image Labeler," <http://images.google.com/imagelabeler/>, last accessed October 26, 2006.
10 Riya Inc., "Riya – Visual Search," <http://www.riya.com/>, last accessed October 26, 2006.
11 Polar Rose, <http://www.polarrose.com/>, last accessed October 26, 2006.
12 Amazon.com Inc., "Amazon Mechanical Turk," <http://mturk.com>, last accessed October 26, 2006.
13 Mycroft Inc., "Mycroft – Technology Needs People," <http://mycroftnetwork.com/>, last accessed October 26, 2006.
14 A nice example of the latter is Aaron Coblin's Sheep Market <http://www.thesheepmarket.com/>.
15 MIT Center for Collective Intelligence, <http://cci.mit.edu/>, last accessed October 26, 2006.

of you shopping. Were this image to be processed by the ESP Game, it is very unlikely that you would be recognized. For this to happen two people, both of whom know what you look like and are able to recognize that the photo depicts you, would have to choose to play the game *and* happen to get matched up as collaborators by the system. But what if the system knew who your friends were and could present your image only to them? Suddenly the possibility of recognition becomes much greater. This is the approach taken by Facebook, a social networking site which rose to prominence due to to its widespread use on college campuses. Facebook allows users to upload photos to their profiles. Friends and contacts of that user can then identify the people depicted in the photos. Pending confirmation by the photo owner, these photos are then linked to the Facebook profiles of the people depicted.

Surveys of Facebook users have shown that they quickly assemble detailed public dossiers on themselves and their friends, often without regard for how easily this information can be harvested (Gross & Acquisti). Through their normal use of the system, Facebook users link hundreds of time-stamped images to detailed profiles of themselves and their friends. Such labor can produce both training data for facial recognition models and fine-grained contextual information about who was where, and when. The latter can be used to significantly improve the performance of (Davis et al.) or completely replace (Naaman et al.) the former.

A new application domain: Homeland security

Hybrid assemblages of people and machines have recently been proposed or implemented as solutions to a number of controversial problems, with serious implications for visual privacy. In June 2006 Governor Rick Perry of Texas announced a plan to spend $5 million on a system that would put web cameras along the state's border with Mexico (Gonzalez & Ratcliffe). The cameras would be equipped with night vision sensors and accessible to anyone with an Internet connection. Would-be vigilantes could monitor the border from the comfort of their homes and call a toll-free number to report suspicious activities. As of October 2006, the cameras were in place but not yet online due to technical complications (Castillo).

Game designers have suggested that a similar system might serve to improve airport screening procedures (Koster, Walsh). Like the designers of the ESP Game, they propose that the problem of airport screening be re-imagined as a massively multi-player game in which distributed crowds compete to identify dangerous people without accidentally triggering searches of innocents. Though the game designers present this idea more as a thought experiment than a serious solution, in principle it does not differ much from Governor Perry's solution to the problem of border security. Should the Transportation Security Authority decide to adopt such an approach, they'll be glad to know that one game company has already done most of the work for them: Persuasive Games. In their game Airport Security players can test their baggage-screening skills from their cell phones or on the web.[16]

Implications for privacy

Recognition links images to people. Increasingly, the time and location at which images were captured is automatically recorded, so that recognition links people to particular times and places as well. Thus recognition technologies carry all the privacy baggage of location-sensing technologies. Like location sensing technologies, recognition technologies can also reveal information about social networks, by

---

16 Persuasive Games LLC, "The Arcade Wire: Airport Security,"
    <http://www.persuasivegames.com/games/game.aspx?game=arcadewireairport>, last accessed October 26, 2006.

revealing who was in the same place at the same time. But recognition technologies, by virtue of being visual, reveal far more than this. Location sensing technologies may tell us that two people were at the same party, but only recognition technologies can tell us whether they were simply in the same room, or whether they had their arms around each other. The image provides another dimension of visual information which we can use for interpreting the situation.

Recognition technologies also make searchable what was not previously considered searchable. I might be comfortable knowing that thousands of pictures of me are scattered around the photo albums and hard drives of my friends and family, maybe including a few here and there that have been made publicly visible on the web. I may be lulled into a false sense of comfort by current image search engines, which rely on image file names or surrounding text and thus only tend to find "official" photos from news stories or biographical profiles. But if companies like Riya and Polar Rose succeed in their stated goals, all of these photos will be collected into a result set with a single query. Careful management of permissions and security in theory can prevent private photos from being publicly viewed, but we know from our experiences with text documents that private information can easily make its way accidentally into search engine indexes. And even if accidents never happened, how is one to manage the complexities of access control? Just the simple act of allowing friend to identify people in photos in Facebook has resulted in a system that collapses social contexts, making one's every appearance in an image available to anyone in one's social network. People are just beginning to take notice of the implications of making personal data searchable and interconnected (St. John), even when access is theoretically limited to one's personal contacts. Now this data includes images as well.

Ethical design and implementation

In the remainder of the paper I present some directions for the ethical design, implementation, and regulation of recognition markets. These include enforcing the contextually appropriate use of metadata, tracking provenance to combat recognition spam, understanding the limits of technological privacy protections, defining the responsibilities of image aggregators, and educating participants in recognition markets.

Enforcing the contextually appropriate use of metadata

Annotations of images, and in particular annotations creating links between images and the people depicted in those images, need to be treated with the same respect for contextual privacy that the images themselves are. If a user of a photo-sharing service creates a link between an image and a person with the sole intent that the link will be used to allow her to easily find that image again, that link should not by default be used to allow global search for images of that person by *any* user of the service. Every effort should be made to ensure that image labelers understand the consequences of their actions, and that the uses to which labels are put do not change after the decision to label has been made.

One might think that this could be achieved simply by allowing the setting of access permissions on images. But even if the image itself is marked as private and not displayed to random strangers, there are ways in which the annotation might be used that violate contextual privacy. Consider for example a display of how many photos in the system depict a given person, shown on that person's public profile. More subtly, consider the use of those annotations to train a statistical model to recognize that person. Thought no one is seeing the images being used as training data, so that the access permissions are

technically being respected, the labels themselves are being used in a way that is inconsistent with the intent of the annotation creator. Just as authors often are unaware that their words will be used as input for marketing algorithms when they write to friends with free web-based email accounts, or post to a site that uses contextual advertising, neither do photo annotators necessarily know that their tags may be used as input for recognition algorithms.

Designers should seek to ensure that participants in recognition markets know for what purposes their annotation labor will be used, and should disallow novel uses of those annotations without the permission of the annotation creators. Imagine that there were a large-scale game in which users competed to find photos which depicted the same person. The effect of such a game is to connect a chain of specific places and times with the the person depicted. Now suppose that, unbeknownst to the players, this information were being used to track political protesters, in some cases resulting in their detainment or arrest. Imagine the psychological damage likely to be suffered were the players to discover that they had been unwitting participants in a reverse Milgram experiment, where what seemed to be just a game actually had quite serious consequences. In the interest of preventing incidents like this, participants in recognition systems need to be fully informed from the beginning about what data is being collected from their activity and how this data will be used.

Tracking provenance to combat recognition spam

As recognition markets grow in size, scope, and economic and strategic importance, there will be increasing efforts to sabotage the recognition process. The creators of the ESP Game recognized the potential for such abuse and outlined some defenses against the most straightforward attacks (Von Ahn & Dabbish). But as organizations which rely on a combination of human-contributed data and machine learning to fuel their economic engines have learned, fighting these abuses is a Sisyphean challenge. Web search and advertising companies spend enormous amounts of time and money fighting search index spam and advertising click fraud, yet it is still unclear whether they are solving these problems or merely keeping them at bay. Recognition markets are likely to face similar issues, with distributed groups coordinating their activity for economic or political gain. For this reason, system designers should be extremely careful to track the provenance of recognition labels so that they can provide audit trails in case of claims of willful mis-recognition.

Understanding the limits of technological privacy protections

Current proposals for privacy-enhancing technologies to defeat recognition systems focus on the exploitation of weaknesses of signal processing-based approaches or the implementation of standards for crippling automated systems. The former includes things like masks or clothing that disrupt recognition algorithms (Alexander & Smith), while the latter focuses on the development of equivalents to "robots.txt,"[17] a file placed on web servers which automated indexing engines are expected to respect lest they be banned from accessing those servers. Neither of these approaches will be very effective against recognition markets. Human recognition is not as susceptible to disruption as machine recognition. Even hiding behind a veil may not interfere with a person's ability to recognize someone based on body shape, gait, or other characteristics. "Do not recognize" flags, on the other hand, are unenforceable against distributed groups. While a photo sharing service may be able to prevent a robotic recognizer that does

---

17 An explanation of the Robots Exclusion Protocol can be found at
   <http://www.robotstxt.org/wc/exclusion.html#robotstxt>.

not respect such flags from crawling its public photos, it cannot effectively prevent millions of users accessing those photos from various places around the world from pooling their recognition ability to achieve the same ends.

Defining the responsibilities of image aggregators

In 2003, the California legislature passed Senate Bill 1386, which requires companies storing personal information about California customers to notify these customers in the event of a security breach resulting in theft of that data. This law created a powerful incentive for companies to invest in securing customer data, rather than simply focusing on keeping such incidents out of the press. But the kinds of detailed personal information which can be produced in a recognition market may not fall under the definition of personal information provided in SB1386. In fact, each individual piece of information may be freely contributed to a public database by users who are unaware of what can be done with their contributions. In recognition markets it is the comprehensive linking of various publicly or semi-publicly available pieces of data which threaten privacy, not the individual pieces themselves. Thus it is worth considering legislation to make visual data aggregators responsible for misuses of the data they collect. This might provide a stronger incentive for companies building recognition markets to ensure that participants are well-informed about how the labels they are contributing will be used, and that these labels will not be aggregated or used as training data in ways that threaten privacy.

Educating participants in recognition markets

Ultimately, the best solution for protecting visual privacy in recognition markets may be education. In particular, people need to be aware of the consequences of indexing images for searchability and of linking images to other kinds of data. In the past people have been prodded into awareness by the sudden appearance of systems that thrust previously submerged issues into the limelight.[18] Designers of these systems sometimes make the argument that the only way to establish reasonable limits for privacy is to keep pushing the envelope until "failure"—that is, until people begin to have negative reactions to the system. This is akin to finding the sharp edges of an object by trying to cut oneself—effective but painful. The danger of this approach is that there may not be sudden, apparent failures until it is far too late to halt the growth and development of recognition markets. Regulating these markets at a late stage of development may prove difficult, if not impossible. Another possibility is that there will never be a shocking failure, but that people will slowly adjust to redefinition of acceptable privacy violation, like the proverbial frog being boiled. In either case, the "build it and see what happens" approach seems unreasonably short-sighted.

Perhaps the most effective brake on the development of recognition markets would be for participants to recognize the value of their labor. The economics of peer production rely on participants under-valuing their contributions, making it possible for savvy organizations to leverage their work for enormous profit. This isn't necessarily bad: as Terranova points out, the participants may be having fun and finding fulfillment in the creative and affective labor in which they are engaging, and they may not particularly care about ownership rights in what they are creating. This is why the commons that Benkler celebrates can exist. But in practice there is often little distinction made between adding to the commons and donating labor to a profit-making organization. If people begin to feel exploited they may begin to

---

18 In addition to the recent Facebook incident described by St. John, the appearance of the first large web archives of old web pages and newsgroup messages in the 1990s prompted waves of concern (e.g. Lasica).

demand a greater piece of the pie, as did AOL chat room volunteers in 1999 (Margonelli), or leave en masse to start competing projects, as did CDDB contributors in 1998 (Lemos). This kind of reaction could make peer production sufficiently expensive or difficult that recognition markets cannot be built.

Some readers may view the suggestions made here as overly critical of peer production or blind to the potential benefits of recognition markets. My view is that peer production is an interesting and potentially quite useful means for achieving some ends, but the ends toward which it is applied are not inherently virtuous. We can damage through open collaboration just as easily (perhaps more easily) than we can build. Others may argue that we are too early in our exploration of the design space of such systems to begin prescribing guidelines. While I agree that we ought not stop experimenting and exploring, it is never too early to begin integrating ethical considerations into our research and design practices. Now is the time for debates about how recognition markets should work and how they will be applied. It will only become more difficult to see how things might have been otherwise.

Alexander, James, and Jonathan Smith. "Engineering Privacy in Public: Confounding Face Recognition." *PET 2003: Proceedings of the 3rd Privacy Enhancing Technologies Workshop.* (2003): 88-106.

Beenen, Gerard, et al. "Using social psychology to motivate contributions to online communities." *CSCW '04: Proceedings of the 2004 ACM conference on Computer supported cooperative work.* (2004): 212-221.

Benkler, Yochai. "Coase's Penguin, or, Linux and "The Nature of the Firm"." *The Yale Law Journal.* 112.3 (2002): 369-446.

Castillo, Juan. "Cameras on border work, but system not on Web." *Austin American-Statesman.* (October 5, 2006): B1.

Davis, Marc, et al. "Towards context-aware face recognition." *MULTIMEDIA '05: Proceedings of the 13th annual ACM international conference on Multimedia.* (2005): 483-486.

Duguid, Paul. "Limits of self–organization: Peer production and 'laws of quality.'" *First Monday.* 11.10 (2006). <http://firstmonday.org/issues/issue11_10/duguid/index.html>.

Frauenfelder, Mark. "Dirty snitches earn $50 for fingering fellow students smoking pot." *Boing Boing.* (April 28, 2006). <http://www.boingboing.net/2006/04/28/dirty_snitches_earn_.html>.

Gonzalez, John W., and R. G. Ratcliffe. "Eyes of Texas to be on border." *Houston Chronicle.* (June 2, 2006): A1.

Gross, Ralph, and Alessandro Acquisti. "Information revelation and privacy in online social networks." *WPES '05: Proceedings of the 2005 ACM workshop on Privacy in the electronic society.* (2005): 71-80.

Hearst, Marti A., Robin D. Hunson, and David. G. Stork. "Building intelligent systems one e-citizen at a time." *IEEE Intelligent Systems and Their Applications.* 14.3 (1999): 16-20.

Kelly, Sean Uberoi, Christopher Sung, and Shelly Farnham. "Designing for improved social responsibility, user participation and content in on-line communities." *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems.* (2002): 391-398.

Koster, Raph. "Treating players like numbers." (January 4, 2006). <http://www.raphkoster.com/?p=242>.

Lasica, Joseph D. "Your Past Is Your Future, Web-Wise." *Washington Post.* (October 11, 1998): C1.

Lemos, Robert. "Companies fight over CD listings, leaving the public behind." *CNET News.com.* (May 24, 2001). <http://news.com.com/Access+denied+A+copyright+battle/2009-1023_3-258109.html>.

Levin, Daniel T. "Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit." *Journal of Experimental Psychology.* 129.4 (2000): 559-574.

Lew, Michael, et al. "Content-based multimedia information retrieval: State of the art and challenges." *ACM Trans. Multimedia Comput. Commun. Appl.*. 2.1 (2006): 1-19.

Ludford, Pamela J., et al. "Think different: increasing online community participation using uniqueness and group dissimilarity." *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*. (2004): 631-638.

Muller, Henning, et al. "Using heterogeneous annotation and visual information for the benchmarking of image retrieval systems." Santini, Simone, Raimondo Schettini, and Theo Gevers, eds. *Internet Imaging VII*. SPIE-6061 (2006): 34-45.

Naaman, Mor, et al. "Leveraging context to resolve identity in photo albums." *JCDL '05: Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*. (2005): 178-187.

Naphade, Milind R., and John R. Smith. "On the detection of semantic concepts at TRECVID." *MULTIMEDIA '04: Proceedings of the 12th annual ACM international conference on Multimedia*. (2004): 660-667.

St. John, Warren. "When Information Becomes T.M.I." *The New York Times*. (September 10, 2006).

Terranova, Tiziana. "Free Labor: Producing Culture for the Digital Economy." *Social Text*. 18.2 (2000): 33-58. <http://muse.jhu.edu/journals/social_text/v018/18.2terranova.html>.

Von Ahn, Luis, Ruoran Liu, and Manuel Blum. "Peekaboom: a game for locating objects in images." *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*. (2006): 55-64.

Von Ahn, Luis and Laura Dabbish. "Labeling images with a computer game." *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*. (2004): 319-326.

Walsh, Tony. "Airport Screening Is A Badly-Designed Game." Clickable Culture. (April 15, 2006). <http://www.secretlair.com/?/clickableculture/entry/airport_screening_is_a_badly_designed_game/>.

Warner, Julian. "Forms of labour in information systems." *Information Research*. 7.4 (2002). 27 October 2006 <http://informationr.net/ir/7-4/paper135.html>.