

Tagging the physical world

Yoelle S. Maarek^{*1}, Natalia Marmasse⁺, Yaakov Navon⁺, Vova Soroka⁺

^{*}Google Inc.

Haifa 31905, ISRAEL
yoellem@gmail.com

⁺IBM Haifa Research Lab
Mount Carmel

Haifa 31905, ISRAEL
{nmarmas, navony, vladi}@il.ibm.com

ABSTRACT

Since Vannevar Bush's futuristic vision in 1945 [1] and one of its more recent incarnations, namely MyLifeBits [4], computer scientists have discussed the benefits of capturing data from the physical world (or rather a "mirror" of them) and bringing them to the virtual world, where they can be transformed into usable knowledge. We propose here to revisit this vision following a social tagging approach, where users intentionally tag objects, people, etc. using a state-of-the-art mobile phone, rather than a dedicated device. In our model, tags correspond to different facets, some automatically captured at tagging time (like location, time, etc) and some manually added like voice/text annotations or discrete values (such as ACL) or even inferred by scanning for instance the images being captured. One key point of our work is our claim that the intentional act of tagging objects, as opposed to recording every bit of information by default, is critical to making the virtual mirror of the physical data useful. The issue at hand here is not storage, which has been shown to be solvable, but the need for users to validate the importance of objects, in a similar way that links carry information on the authority of pages in the Web. Some directions for future research in a field, which we believe is only in its infancy, are also sketched out.

Keywords

Social tagging, Web 2.0, social networks.

1.INTRODUCTION

With the progress recently made in storing, searching and retrieving all kinds of data, from Web to personal desktop information, in so many different formats, Internet-savvy users become sometimes frustrated with the limitations of the real physical world. Haven't we all once wished, while shopping in a large appliance store, that we could have immediate access to online opinions about this new LCD large screen? Wouldn't it be nice to know whether the price is really competitive, how it compares to plasma screens while looking at it in the store? Some innovative retail stores such as Walmart have already started experimenting with RFID technology for inventory purposes [14]. This can be seen as a first step in bridging the gap between the physical and virtual worlds.

Another frustration that we have all experienced is our inability to retrieve memories of places, names, and facts in the same way that we can search our desktop. Even if one of the first desktop applications was called "Stuffed I've Seen" [3], it referred to "stuff" seen in the virtual world and not in the real world. An

interesting and much talked about experiment, designed to precisely address this need, is MyLifeBits [4], which implements the visionary "Memex" device discussed by Vannevar Bush in 1945 [1]. Impressively enough Bush envisioned not only the ability to search the individual and global knowledge but maybe even collaborative tagging, through the ability to annotate information and establish "trails" (read links). Gordon Bell has been running this experiment on himself for the last few years and claims to have "become paperless since 2002": he has scanned and stored books, personal documents, CD, photos etc. but more importantly he has started to record phone conversations, radio broadcasts, TV shows, meetings, etc. Bell also investigated the benefits of annotations on pictures, etc. and how it will help automatically build correlations. Figure 1 shows an illustration of Vannevar Bush's Memex device side by side with one of Gordon Bell's recording devices. As last reported in 2004, Bell's "MyLifeBits" has taken only 30GB of storage space, and it was demonstrated that with the progress in storage, keeping one's personal information is clearly achievable in terms of both cost and size.




Figure 1: Vannevar Bush's vision and Gordon Bell's implementation

Another project inspired by the Memex is Life Patterns [2], which continuously captures and stores audio and video from wearable sensors. Using computational models, it extracts the recurring events of daily life; these are the raw material of our memories. As Clarkson puts it "memories are a sort of compression of our physical experiences". StartleCam [5], on the other hand, aimed to record non-routine experiences. The system included a wearable computer, camera and galvanic skin response (GSR) sensor – which provides measurements associated with stress or arousal, such as created by a startle response. The GSR levels are monitored and when a startle is detected, a picture of the

¹ This work was conducted while the author was still with IBM Haifa Research Lab

“significant” moment is taken. The “What was I Thinking” memory prosthesis [15] indexes continuously recorded audio with “memory triggers” e.g. location, weather, related email and calendar entries. Yet another Memex research prototype, which has been used/worn for several years, is the Remembrance Agent [11]. While such projects are a major breakthrough in realizing the Memex vision, they are still far from being practical. As Gordon Bell puts it, he is a “guinea” pig and we don’t envision many followers yet ready to scan manually every piece of data they own. Worse, we cannot imagine many “normal people” going around with a camera recording everything they see, hear or say. We believe that one major show stopper for Vannevar Bush’s vision to be realized is the user experience.

One of the best user experiences that we have seen so far in this area is unfortunately only imaginary and even “magical” ☺, as it is taken from *Harry Potter and the Goblet of Fire* (4th volume of the series) [12], where the readers discover Dumbledore’s *Pensieve*, a magical stone basin that stores memories and thoughts. The *Pensieve* even takes a more important place in the last volume of the series, *Harry Potter and the Half Blood Prince* as it made the cover in some editions, see Figure 2.



What is it?" Harry asked shakily. "This? It is called a Pensieve," said Dumbledore. "I sometimes find, and I am sure you know the feeling, that I simply have too many thoughts and memories crammed into my mind." [...] "At these times," said Dumbledore, indicating the stone basin, "I use the Pensieve. One simply siphons the excess thoughts from one's leisure. It becomes easier to spot patterns and links, you understand, when they are in this form [...]"

"Harry Potter and the Goblet of Fire", chapter 37

Figure 2: Dumbledore’s Pensieve

One major innovation in the *Pensieve* is that memories and thoughts are committed to the stone basin by a “simple” action: Dumbledore fetches them by dragging² them from his head and dropping them into the basin with the help of his wand. The device, the wand, always resides with its owner and is not used only for operating the *Pensieve*. More importantly, the action of committing memories to the *Pensieve* is *intentional*.

We see this act as extremely reminiscent of social tagging. Social tagging systems, which allow users to select specific Web pages, such as www.furl.com, del.icio.us, www.rawsugar.com, www.simpfy.com, etc., Intranet pages such as in [9], as well as arbitrary objects such as images, www.flickr.com are by essence intentional, users actively select a page or objects they want to keep or share. Similarly, Dumbledore actively tagged³ his memories and stored them for later retrieval as well as for spotting

² The “drag and drop” parallel is intentional here.

³ The similarity with social tagging stops here though as the *Pensieve* does not support annotations allows only a limited form of sharing – Dumbledore can only show his memories to Harry and not draw correlations between common memories

“patterns and links”, to cite the book again, See Figure 2. One additional critical point is that by committing memories to the *Pensieve*, Dumbledore implicitly gives them some degree of importance. He would not commit to the *Pensieve* for his instance his memory of having brushed his teeth in the morning. This is critical to future usage and to the future retrieval in particular, in the same way that manually establishing links between Web pages confer some of these pages authority over others. This linkage is what allows modern search engines to boost their pure content analysis technique so as to identify authoritative pages on a certain topic [7].

Returning to our real world, we propose here to draw a parallel between memories committed to the *Pensieve* and snapshots of the physical world surrounding an individual brought to a personal virtual repository. Since we cannot capture and model memories *per se*, we propose a significantly simpler approach, where users, instead of tagging a memory with a wand, capture an image or an audio record of an object, or a person that they are exposed to at a given moment. That is, capture some of the subcomponents of which memories are formed.

We discuss in the rest of this paper how the capturing act is carried out, through a readily available state-of-the-art mobile phone. Section 2 presents our approach, introduces the tagging model for representing these pseudo memories in the virtual world, around a faceted scheme. Section 3 discusses some usage scenarios and present an actual case study carried with our first prototype. We conclude with directions for future research. Related work is discussed as relevant throughout the paper.

2. OUR APPROACH

2.1 Capturing memories

Since telepathic capabilities are obviously beyond our reach, we cannot hope to implement *Pensieve*’s function and capture thoughts and memories. We therefore set ourselves a more modest goal: namely to capture memories derived from the physical world. These actual memories are represented via a surrogate in the virtual world. This surrogate can be seen as a mirror of the memory and will immediately trigger the actual memory upon retrieval. Memories can include actual elements of the world such as: Persons (faces, gestures, positions, etc.), nature (flowers, birds, scenic views, etc.) or monuments (buildings, bridges, etc.).

They can relate to primary objects as listed above, but also to secondary elements that can optionally be linked to a primary element. Examples of such secondary elements can be a business card, which relate to person, a page from a book that relates to a topic, a poster of a conference, relating to an event, an advertisement (roadside billboard, or in a journal) that relates to a given product or service, etc.

The common denominator between all the above examples is that they are all in printed form and can then be OCRed and processed so as to automatically extract attributes. Thus business cards will give us the attributes of a person (as detailed later), posters will give us attributes of a conference. As an example the WWW’2006 poster of flyer seen on the wall, could be photographed, and OCRed to extract its URL or location.

The images and videos can naturally also be annotated with audio via voice recording. Take as an example, Jane meeting a new colleague named Joel, she will be able to capture the memory of

Joel via a picture she would take, getting his business card as explained above, she could at the same time pronounce and save his name or even record Joel saying his own name because he is from France and pronounces it differently. In addition to location, time and date the images may also include extended personal data (name, phone, car model, etc.) of the mobile's owner. Tagging like "I took the Eiffel's picture when I owned a Buick 2006..." is then feasible.

While we distinguished here between primary objects, secondary objects and even attributes, we argue that a universal model of the world where a limited set of semantic objects need be defined is inappropriate. We here fully adopt the more recent tagging models inferred by del.icio.us, flickr, Rawsugar, etc. where tags emerge from the community and semantic links are automatically built by users. We therefore allow a totally open uncontrolled vocabulary of tags, and like in modern social tagging systems, the objects will stabilize once enough objects are tagged. We refine this approach in 2.3.

2.2 The capturing device

Our wand is a device ubiquitous among *muggles*⁴, namely a Java-enabled mobile phone equipped with a camera. In a study on the ways in which people use their camera phones [6], it was found that two-thirds of images taken for functional reasons were to support a practical task, e.g. recording information for future reference. Many images are also taken for posterior reflection or reminiscing. As the study points out, there is a need for better tools to help users sift through the captured images.

We have experimented with mobile phones that support Java 2 Micro-Edition (J2ME), with the Connected Limited Device Configuration 1.1 (CLDC) and Mobile Information Device Profile (MIDP) 2.0 APIs. These include the Nokia 6680 with its 1.3 megapixel camera and the Sony-Ericsson W800 with a 2.0 megapixel camera.

A Java application on the phone enables a user to easily snap photos of places and/or people he wishes to remember, as well as record an audio snippet either annotating the moment or simply recording the memorable ambient sounds. The captured images and audio are then transferred via Bluetooth or GPRS to the user's desktop or to a backend server, where the images may be processed and inserted in a calendar and/or address book application. Optionally a Bluetooth-enabled GPS can be used to enhance the captured memories with their geographic location.

2.3 A faceted model

In our model, we associate "captured memories" with facets at different levels of granularity. Two of the highest level facets, given our tagging device, are the **image** facet, whose value is the actual JPEG file and the **audio** facet whose value is the audio record itself. These facets can be refined by qualification e.g., in "**image>format:JPEG**", the **format** facet takes JPEG as value and refines **image**.

We do not enforce a strict hierarchy among facets and use them the same way they are used in social tagging systems. They can relate to each other or not. Our facets can be built bottom-up, top-down, reorganized after the fact, etc. They are primary objects on which fielded (or faceted) search can be applied.

We distinguish between three types of facets:

- **Implicit facets:** These facets are automatically instantiated by the tagging device. The example devices described in the previous section will automatically instantiate the **date/time** facet for instance. However, much finer grain information can be expected now or in the near future. Indeed, most modern cameras store images in JPEG augmented with EXIF⁵ (Exchangeable Image File), which account not only for camera settings (see flickr data for instance) but also location information for the rare devices that do already have GPS capabilities
- **Explicit facets:** These facets are manually added by the user or prompted for by the tagging application. Thus, if the user often captures the same type of object, such as a business card, s/he might decide to define this as a facet and select it from a menu at tagging time, in order to ease further processing and/or retrieval. Users can also, in advanced devices with input capability simply enter free-format tags, or alternatively enter them after the fact in the associated client application. An audio description of an image (as opposed to an audio "memory") can be seen as an explicit facet decorating an image memory (e.g., "This picture shows Joel graduating").
- **Derived facets:** These facets are the most intriguing as they are automatically derived from data. Thus, all the information that can automatically be extracted from an image via OCR for instance, or from an audio record via ASR technologies, fall under this category. Our case study in the following section, studies more in depth the facets that can be derived from specific types of images.

2.4 Sharing memories

While we have mostly focused until now on the personal aspects of the tagging exercise, like in social tagging, most of the value will be inferred in the long term, from communities of taggers. A communal repository storing the "memories" of entire communities should bring incredible value. Though we are still far from attaining this goal, we are convinced that the communal repository of memories will allow us to fix many of the uncertainties, to augment everyone's knowledge and discover unexpected patterns

Imagine that several individuals take a picture of a same poster, of a same person, of the same business card, and that enough information has been inferred or derived so as to guarantee that all these snapshots correspond to the same "memory". Simply having at our disposal various pictures taken from different angles, under different lighting conditions, by different devices will in most cases allow us to significantly improve the quality of OCR as a trivial example. In addition, tagging the memory of a same person with different facets will give us a more general surrogate of that same person. Basically, as immediate benefits of collaborative tagging, we expect significant improvements in both precision (the quality of the derived facets) and in recall (the coverage of facets for a given memory). Finally, one can also envision that mining huge communal repositories will bring us to new levels of discovery.

⁴ Non-magical folk, in Harry Potter's world.

⁵ <http://en.wikipedia.org/wiki/Exif>

3.CASE STUDY AND PRELIMINARY RESULTS

Tagging of the real-world has numerous uses, and one can foresee new ones emerging as technology becomes more prevalent. We enumerate here just a few examples, most of which are relevant both for personal and communal repositories. Repositories may typically be shared between small groups which are emotionally close-knit (e.g. family and friends) or task-related (e.g. colleagues at work), or even more globally.

3.1 Usage examples

Imagine the following scenario where the user tags “memories” of people or of artifacts related to people s/he meets so as to populate and enrich his/her contact organizer application. The user, via the camera phone, takes a collection of photos, such as, one of the person he has just met, of his business card, the conference booth where they met, etc. The user may also record snippets of their conversation, summarize the meeting, or add action items in a voice annotation. The contact organizer application, using OCR, extracts the derived facets such as the person's contact information (name, email address, etc), creates an address book entry for this contact, and associates all images, audio, and any of the implicit or inferred context information to the new contact. If the address book already contains the contact, the new collection is added to the existing ones. Similarly, the new “memory” components can be linked to a Calendar entry. Searching, either by photographing a business card or manually inputting a name or email address, would retrieve all the previously recorded memories/encounters of this person.

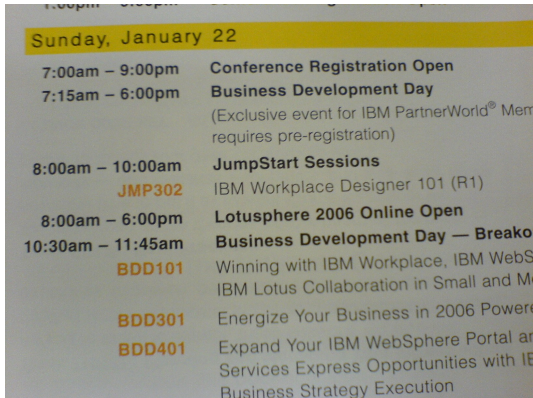


Figure 3: Snapshot of the pocket agenda of LotusSphere 2006

The camera phone can also be used to collect information on future events from posters or newspapers. It would then be possible to extract the relevant information, such as the event time and date, URL, etc., and associates them with a new calendar entry. For example, the image depicted in Figure 3 would automatically be entered in a Calendar entry for January 22nd. Likewise, images or audio serving as reminders, whether for personal use or directed for someone else, can be associated with a calendar entry or to-do list. One can also envision the capturing/tagging and broad sharing of amusing billboards, strange optical illusions, sound effects, and so forth.

3.2 Our case study: Attending a conference

We decided to implement a prototype of the first application envisioned in Section 3.1 for organizing personal contacts and experiment with it while attending a real conference. We called

this early prototype Pensieve as a credit to Dumbledore's [12]. We used a Sony Erickson w800i as our tagging device, and reused in-house OCR technology. Figure 4 shows the image of a business card captured by our phone. The connected components analysis process results in two text blocks, which are then binarized and enhanced for OCR processing.



Figure 4: A business card captured by a mobile phone camera

The actual raw OCR results and the fields extracted by the disambiguation process are shown in Figure 5.




Figure 5: OCR results of business card

The tag processing component of our Pensieve conducts some of its tasks locally right after capturing and some after the fact, on the user's device depending on the computational burden. Most of the OCR and post OCR analysis for instance is done on the user's desktop. We distinguish here between the tags that are derived from the meta-data of the picture, those which are manually given by the user and finally those which are automatically derived after the OCR and type-sensitive disambiguation process. Facets are indicated in bold font and their value follows, after the semi-colon character.

The *implicit* tags taken from the image meta-data are:

- **Image_name:** DSC00083.JPG
- **Date:** 06:12:2005
- **Time:** 18:33 GMT+2

The only manually entered *explicit* tag is

- **Type:** bcard.

Once the image type is known, some semantic rules are applied that take advantage of the specific format. In our example case, knowing the type allows adding these *derived* tags:

- **Email:** navony@il.ibm.com
- **Name:** Yaakov Navon PhD
- **First name:** Yaakov
- **Last name:** Navon
- **Title:** PhD
- **Body:** Image Technologies Depadment
IBM Research Lab
- **Tel:** + 97248296132

- **Mobile: +97253790486**
- **Fax: +97248296112**

Note that further correction and finer tags can be derived if required. For instance simple spelling correction can turn *Depadment* into *Department*.

This information is used to populate the Address book in our prototype. The time of the picture is used to retrieve the calendar information to define the event associated with the business card.

3.3 Preliminary results

Visiting a conference is one of the real life settings where meeting people is one of the primary activities. One of the authors attended Lotusphere 2006 [8] and took advantage of his sitting at the conference Innovation Lab to experiment with our Pensieve prototype. He systematically documented his meetings with 112 people by tagging them with Pensieve. The main goal of this preliminary study was to assess the usefulness of the concept and record potential user's reaction to it.

Of the 112 pictures taken, 77% were considered useful after OCR. The remaining 23% were unclassified (OCR failures), the primary reason being blurred pictures; the camera was also, in a few cases, understandably sensitive to fancy business card design, were the unusually glossy paper created reflections or where the name was written in a very unusual font. In the majority of the cases, our prototype succeeded in recognizing the name of the contact and adding it to the address book. In 10% of the cases, the OCR was perfect, collecting exhaustively all the information listed on the business card.

Despite the limitations of our OCR component, which we do plan to tune for better results in the very near future, the people we interviewed were very impressed with the concept, especially when Pensieve got their names right. Many people stated that "the tool is perfect for salesmen, meeting a lot of people daily". One of the visitors even said that in her company employees are required to document their encounters with people after returning from any event, using a special database application. The "one-click" concept—obtaining all the information available in one click—*on-site* also impressed a lot of potential users, which found it clearly preferable to available commercial solutions, where users use a special dedicated scanner typically upon returning from their trip and process all their accumulated business cards in batch mode.

Finally, the real potential value of Pensieve can be inferred from the following anecdote. After the conference, when going through the profiles of captured business cards, the demoer found an entry of a lady he used to collaborate with via email several months before. He actually really wanted to meet her personally. If our Pensieve implementation was already tied to one's social network, he could have recognized this on the spot when "tagging her" and could have mentioned their previous interaction. These kinds of "on the spot", live connections are just an example of how social networks can be brought to the physical world.

4. CHALLENGES

4.1 Technical – OCR

The main input to our current Pensieve prototype are images, hence we limit our technical discussion to OCR (Optical Character Recognition). Future versions may include the transcription of recorded voice snippets, leading to additional, audio-related, challenges.

Our OCR process includes several steps, such as binarization and connected component-based layout analysis, before character recognition can be applied to the different blocks of the binary image. The final stage is the disambiguation process, where the extracted text is corrected and provided with the relevant descriptors. In the business card example, these are: name, email, telephone, company name, position, etc. The various OCR sub-processes become significantly more complex when carried out on images captured by a camera phone, as opposed to images scanned by a well calibrated flat-bed scanner. Moreover, camera phones generate compressed lossy images (JPEG) which suffer from "blocking artifacts"[10], and bad reconstruction of image areas where sharp contrasts exist, as in text areas.

Besides the limitations of the camera itself, one must consider the conditions in which the images are taken, adding yet another layer of complexity. The camera phone images are typically shot in varying and difficult lighting conditions, with unstable hands, with perspective distortions, and often by unskilled users who do not necessarily choose the best camera settings.

Our present implementation focuses on OCR of the embedded text, however future versions may include 1D and 2D Barcode decoding, graphic recognition of logos and symbols, face detection and recognition, image analysis, and object recognition (e.g. monuments, building, cars, roads, etc.) based on geographical location data and field of view. Each of these tasks will add further challenges.

4.2 Mobile devices

The first mobile phone to include a digital camera was a Sharp model, introduced in Japan in November 2000. Nowadays, 47% of phones sold in North America have a built-in camera, in Europe it is 55%, and in Japan the figure is up to almost 92% [13]. While today's state of the art camera phones have advanced features, such as auto-focus and zoom, not all these features are available to developers to control via their Java midlets.

4.3 Social

The ability to easily capture data—whether it consists of images, audio, or video—inevitably raises many privacy concerns. What can be captured, when and for what purpose varies from country to country, and in the U.S. even from state to state. For example, in some places, audio can only be recorded if there is consent of all parties present. With mobile phones becoming so popular, many worry about camera phone voyeurism. What legal rights do we have against having our picture taken in public? Some would argue that this falls under the "right to be left alone", articulated in 1890 by Warren and Brandeis in their famous article "The Right to Privacy" [16]. A big concern with digital data is that once it has been saved, a party involved often has no control over who has access to it, or how it will be used. Laws are being amended to reflect the advances in technology, however these laws are often hard to enforce. For example, camera phones are being barred from gyms in Australia and in Japan there is a new crime called "digital shoplifting", where magazine pages are photographed instead of bought.

Perhaps more interesting than changes in the relevant laws, is the impact on our social behavior. How will social mores evolve to incorporate this new technology? Let's take for example the business card. The Japanese traditional greeting in the form of a bow was extended to include the modern business custom of

exchanging cards. Although a relatively newer custom, it follows the ceremonious style of older ones. After the bow, the business cards are given and received with both hands, and then examined carefully as a show of respect. Once business cards become frequently photographed, instead of distributed, we expect the card itself to change, perhaps acquiring a more personalized format. What customs will gradually emerge around the photographing of the card and its owner? How acceptable will it be to record snippets of conversation?

5. CONCLUSION AND FUTURE WORK

This paper presents our vision of tagging the real world, and a faceted-model of the captured "memories". We describe an initial implementation as well as some results and user feedback, which although preliminary lead us to be optimistic both regarding the approach and the desire for such a tool.

In the short-term we plan to improve the OCR rates on images with text, and add recognition of other image types. We would also like to add audio transcription for tagging. Additionally we plan to explore how real-world tagging can benefit communities of users. Similar to *delicio.us*, where people benefit their community by simply maintaining their own bookmarks, we believe that in our scenario people can benefit the community by sharing their tags or "memories". In this case however, issues of privacy will be much more significant, since in the real world there are much more personal things than on the Web. Nevertheless we are confident that an easy and adaptive solution to this problem can be found – this is the main lesson of the Social Web.

6. ACKNOWLEDGMENTS

We are grateful to Stu Feldman, who helped us focus our preliminary ideas. We also thank Ella Barkan who contributed significantly to our OCR technology.

7. REFERENCES

- [1] Bush, V. *As We May Think*, The Atlantic Monthly, July 1945. Available at <http://ccat.sas.upenn.edu/~jod/texts/vannevar.bush.html>
- [2] Clarkson, B. (2002) *Life Patterns: Structure from Wearable Sensors*. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, Sept. 19, 2002.
- [3] Dumais, S., Cuttrel, E., Cadiz, J., Jancke, G., Sarin, R., & Robbins, D. (2003). *Stuff I've Seen: A System for Personal Information Retrieval and Re-Use*. In Proceedings of SIGIR'2003, Toronto, CA, August 2003.
- [4] Gemmell, J., Bell, G. and Lueder, R. *MyLifeBits: A Personal Database of Everything*. Communications of the ACM, 49:01, January 2006
- [5] Healey, J. and Picard, R. W. (1998). *StartleCam: A Cybernetic Wearable Camera*. In Proceedings of the 2nd IEEE International Symposium on Wearable Computers, October 1998.
- [6] Kindberg, T., Spasojevic, M., Fleck, R. and Sellen, A. (2005). *The ubiquitous camera: An in-depth study of camera phone use*. IEEE Pervasive Computing, Vol. 4(2), pp. 42-50.
- [7] Kleinberg J., *Authoritative sources in a hyperlinked environment*. Proc. 9th ACM-SIAM Symposium on Discrete Algorithms, 1998. Extended version in Journal of the ACM 46(1999).
- [8] Lotusphere 2006 conference. <http://www-142.ibm.com/software/sw-lotus/events/govfor.nsf/wdocs/ls2006home>
- [9] Millen, D. R., Feinberg, J., and Kerr, B. *Dogear: Social Bookmarking in the Enterprise*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Montreal, Canada, April 24 - 27, 2006). CHI '06.
- [10] Penebaker, W.B. Joan L. Mitchell, J.L., "JPEG still Image Data Compression Standard", p. 38, Van Nostrand Reinhold, New York, 1993.
- [11] Rhodes, B and Maes, P. (2000) *Just-in-time information retrieval agents*. In IBM Systems Journal special issue on the MIT Media Laboratory, 39, Nos 3 and 4, pp 685—704.
- [12] Rowling, J.K., *Harry Potter and the Goblet of Fire*
- [13] Sharma, D. "Camera phones sales on upswing", CNET Asia reviews, Dec 6, 2005. <http://asia.cnet.com/reviews/handphones/0,39001719,39295802,00.htm>
- [14] Sullivan, L. "Wal-Mart Tests RFID Data-Sharing Project", Information Week, Jan 17, 2006. <http://www.informationweek.com/industries/showArticle.jhtml;jsessionid=PIUSOBFBRBJZH0QSNDBECKH0CJUMKJVN?articleID=177101277&pgno=1>
- [15] Vemuri, S. and Bender, W. (2004) *Next-generation personal memory aids*. In BT Technology Journal. 22(4) 125-138, October 2004.
- [16] Warren, S. and Brandeis, L. *The Right to Privacy*, Harvard Law Review IV:5, Dec 1890. Available at http://www.lawrence.edu/fast/boardmaw/Privacy_brand_war2.html