

A Low-Power, Distributed, Pervasive Healthcare System for Supporting Memory

Dirk de Jager, Alex L. Wood, Geoff V. Merrett, Bashir M. Al-Hashimi,
Kieron O'Hara, Nigel R. Shadbolt and Wendy Hall

{ddj07r, alw306, gvm, bmah, kmo, nrs, wh}@ecs.soton.ac.uk

School of Electronics and Computer Science
University of Southampton, Southampton, SO17 1BJ, UK

ABSTRACT

With age, injury, or disease, human memory can be significantly impaired. Current coping strategies range from simple aids such as post-it notes and calendars to, more recently, assistive devices which attempt to provide reminders at appropriate times, capture details of important events, or aid with performing complex tasks. In this paper we extend the state of the art by reporting on a pervasive system designed to ‘orientate’ a person with memory problems by reminding them in real time of details and personal memories of their surroundings. We describe a three-level architecture, consisting of a novel wearable sensing device, mobile phone, and Internet service, where each level is less constrained in terms of energy and offers greater computational capability. This kind of distributed architecture for a memory aid has not been reported before. The wearable device typically draws 38.5mA of current, allowing a full day’s operation on a single charge. We demonstrate the operation of the system, which successfully identifies the faces of those nearby and shows their names on the mobile phone.

Categories and Subject Descriptors

H.4.0 [Information Systems Applications]: General; J.3 [Life and Medical Sciences]: Health; J.4 [Social and Behavioral Sciences]: Psychology

General Terms

Design, Experimentation, Human Factors

Keywords

Electronic memory aid, assistive technology, low-power mobile device, distributed processing, pervasive healthcare

1. INTRODUCTION

Human memory is not perfect. Dementia, injury, and the natural decline of mental abilities with age all further affect our ability to remember our past, our future plans (whether for the next minute or next year) and the important facts of our lives [1]. Memory impairment has a huge impact on quality of life, and with an ageing western population, it is becoming ever more of a problem [2].

Over the past few decades, a number of attempts have been made to aid memory using technology, with many positive results. A variety of aids have been reported for people who have problems remembering future intentions, such as using automated SMSs and pager messages [3]. Other work has aimed to improve ‘memory of the past’, replacing or augmenting solutions in widespread use such as diaries and photo albums [4].

The classification of memory aids as ‘retrospective’ (past experiences) or ‘prospective’ (future plans) is one based on widely-held psychological theories, in which memory is divided into various functional systems. Very broadly, retrospective memory is divided into episodic and semantic memory. Semantic memory stores general knowledge, while episodic memory is that which records vivid, personal events. Prospective memory relies on components of retrospective memory to record the detail of the task to be done, and also makes use of various other systems for recalling the intention at an appropriate moment [5].

In designing a memory aid, it is essential to focus on providing useful cues to the user (people, places, objects, and actions) to stimulate their own memory [6], designing around psychological models to support the user’s remaining capacities rather than attempting to completely replace their own memory.

The earliest reported work on an intelligent electronic

memory aid was by Lamming & Flynn in 1994 [7]. Observing that workplace efficiency was reduced by employees' inability to remember the details of files, people, and appointments, they proposed a pervasive system as a solution. Interactions with people (wearing IR badges), access of documents on computers, and phone calls (via the internal exchange logs) were recorded, after which data were accessed via a PDA. Contemporary work considered the possibility of offering context-related information in real time via a worn screen [8]. More recently, Cooks' Collage [9] demonstrated the feasibility of using computer vision to aid with a memory task in real time, reporting a system which improved the following of recipes whilst simultaneously performing another task. Mihailidis, et al. developed a similar system which employed computer vision to monitor demented residents' hand-washing, offering step-by-step prompts as necessary [10].

SenseCam [11] was developed as a life-logging device with a focus on photos, consisting of a small wearable camera coupled with desktop photo-viewing software. It used a combination of sensors and a timer to decide on useful times to take a photo. It became quickly apparent that, by means of repeated review of the photos, SenseCam could be a useful aid for improving the quality of life of patients with impaired episodic memory and, more recently, semantic memory [12]. The concept was extended by Lee and Dey [13], who built a system specifically for use by patients with memory impairment. A SenseCam was used to capture photos, and audio was simultaneously recorded via a dictaphone. Improved review software allowed a carer to pull out important cues and have them displayed progressively during review, giving the patient time to remember as much as possible independently before being cued further.

Fenwick, et al. recently outlined a mobile phone-based system which helps with recalling names [14]. Noting a strong structure to elderly peoples' schedules, a system was devised to provide a shortlist of people likely to be encountered in a specific location, along with associated textual aide-memoirs added by the user.

This section has highlighted that, while considerable progress has been made, there is still a gap in capabilities of the various systems in the literature, and a user-friendly, interactive aid for retrospective memory is still not available. For a more extensive review of the state of the art in pervasive technology and memory aids for supporting people with memory problems, see [15] and [16].

In the remainder of this paper, we present the 'Deja-

View' system, which brings together the retrospective and prospective capabilities of existing devices, and additionally introduces real-time feedback such as the last time they met someone currently nearby, or a reminder from their calendar on where they should be going. The design takes advantage of the considerable connectivity offered by modern mobile phones to facilitate intensive computation work in a system with tight energy constraints. The DejaView architecture is validated and evaluated with near-real-time face recognition in an initial prototype, and the energy consumption of the wearable device is shown to be similar to that of the SenseCam. We argue that the DejaView system offers a step change in the capability of electronic memory aids. Specifically, the novel features of the system are:

- A pervasive system which provides memory prompts in real time to help with both orientation in real time and with retrospective memory;
- A low-power, wearable device which intelligently takes photos and transmits them wirelessly to a mobile phone;
- A three-tier distributed architecture allowing intensive processing while observing energy constraints in the portable parts of the system.

The design of the system is described in Section 2. Its performance is validated and evaluated in Section 3, and conclusions and future work are considered in Section 4.

2. THE PROPOSED SYSTEM: DEJAVIEW

The DejaView system is designed to aid recall of daily activities, plans, people, places, and objects by sensing the user's surroundings and inferring context. The system then unobtrusively cues a user with relevant information, helping them orientate themselves and aiding both their prospective and retrospective memory. We consider prospective memory support as provision of useful cues rather a simple alarm. An example of feedback from the system would be: 'you are speaking to Anne, who you last saw on Wednesday afternoon at Lily's with Roger and Colin. Check she remembers about lunch with Mary tomorrow at one o'clock in the park'.

Facilitated by a three-tier architecture, the system analyses context with computationally-intensive processing while observing strict energy requirements for the wearable device (Figure 1). This architecture consists of:

- A low-power wearable sensing device, called the

DejaView device (Figure 2), which autonomously takes photos based on inputs to its sensors;

- An application running on an Internet-connected mobile phone;
- An Internet service, running intensive processing such as face recognition, and connecting to the user’s social networks, calendar, online photo albums, etc.

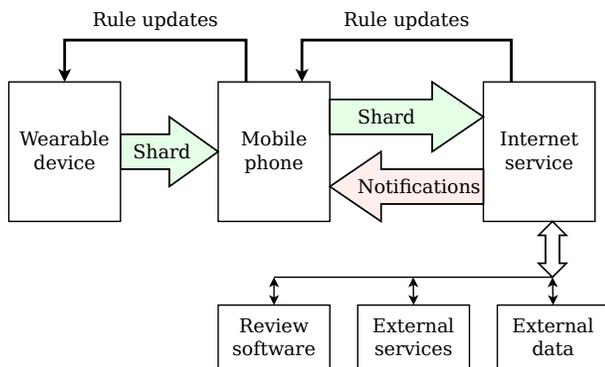


Figure 1: DejaView architecture, showing data flow through the system

2.1 System Overview

The wearable DejaView device (Section 2.2) contains a suite of sensors which constantly monitor the user’s environment. When the sensed values match particular rules, an onboard camera is triggered to capture a photo. This allows the camera, a high-power device, to be turned off until the device has decided to take a photo. The system architecture is summarised in Figure 1.

Once an image has been captured, it is transmitted wirelessly to a mobile phone (Section 2.3) along with the sensor readings and rules which were used to take the image. This is encapsulated in an (*information*) *shard*. The mobile phone then further processes the information, appends additional contextual data, and uploads the shard to an online application for further analysis.

In the Internet service (Section 2.4), the incoming shard is related to the users online presence (checking their online calendar to ensure that they are not far from an existing appointment, or comparing the faces of people in the captured photo with those on their social networks or in past photos). Information extracted and evaluated by this process is then compared to a table of ‘notification rules’. Information about the context (such as peoples’ names, or the time they were last in



Figure 2: The DejaView device (right) and mobile phone (left) evaluated in this study

the same place) is communicated to the user according to these rules. Table 1 gives some examples of the kind of feedback. Rule updates to the DejaView device and the phone are also sent, adjusting sensor thresholds and photo capture rate as appropriate for the current setting to balance latency with power consumption. The photos are saved for later user review, and for future system use where the same context is encountered again.

The information shard attempts to describe a self-contained event based on the sensors’ data. A shard contains the following data or part thereof:

- The captured photo;
- The sensor data from the time of capture;
- The rules which have passed in the various levels.

The information contained in the shard thus allows further stages in the system to understand which rules were used to trigger the capture and the data which triggered the rules, and hence make changes to those rules as desired.

Using a mobile phone as the user interface to the DejaView system removes the requirement for the new device to have its own user interface, decreasing the overall energy requirements of the DejaView device and simplifying the user experience. The usability of the system will nevertheless affect its application: patients with advanced-stage dementia would be unlikely to be able to operate a mobile phone, but those with earlier-stage dementia or other forms of amnesia will be unlikely to have any serious problems with some training.

Table 1: Examples of photos captured by the device, with significant features and possible system responses described

Image	Recognised	Information and response
	Faces: Ian, Steve, Richard	Talking to Ian, Steve, and Richard. Remind the user of when and where they last spoke to these people Important situation. Increase the image capture rate.
	Place: hallway, 34 Greene Road	Walking down hallway. Appointment at 4:30pm. Reminding user of an appointment and where they are. Unimportant situation. Decrease the image capture rate.
	Objects: milk, microwave,	In the kitchen. Has the user left milk out of the fridge? Have they left the fridge open? Important situation. An action may be performed.

The distributed architecture of the system maintains computational power whilst also minimising energy requirements in the wearable and portable parts of the system. The key challenge of the distributed approach is latency: can the user be given information in a useful timeframe? If the user is depending on the system for basic details such as peoples’ names and other basic information, it will need to provide feedback within a few seconds. However, if the user only needs support in remembering more detail and recent history, information being returned within a few tens of seconds could be acceptable. We evaluate the latency of the system in Section 3.2. In addition to the real-time capabilities, it is anticipated that most of the existing software for SenseCam image libraries will also be compatible with DejaView photos, and hence the system will be at least as capable as the SenseCam in supporting retrospective memory through a repeated review process.

2.2 Device

The wearable DejaView device autonomously takes photos based on inputs to its sensors. It must have a battery life approaching a day, and so needs to be low power. The main components of the wearable DejaView device are a low-power microcontroller (MCU), a suite of sensors, low power static memory (SRAM), a camera, and a Bluetooth radio, as shown in Figure 3. The device is powered by a lithium polymer battery, charged through a micro-USB connector. A dedicated battery monitor manages the voltage regulators and power rails for the

microcontroller and sensors. The microcontroller manages the power rails of the camera and radio, turning them off to enter extremely low power modes. Sensed data is captured, processed, and stored to the mobile phone via Bluetooth.

A wide range of sensors were evaluated for use in the

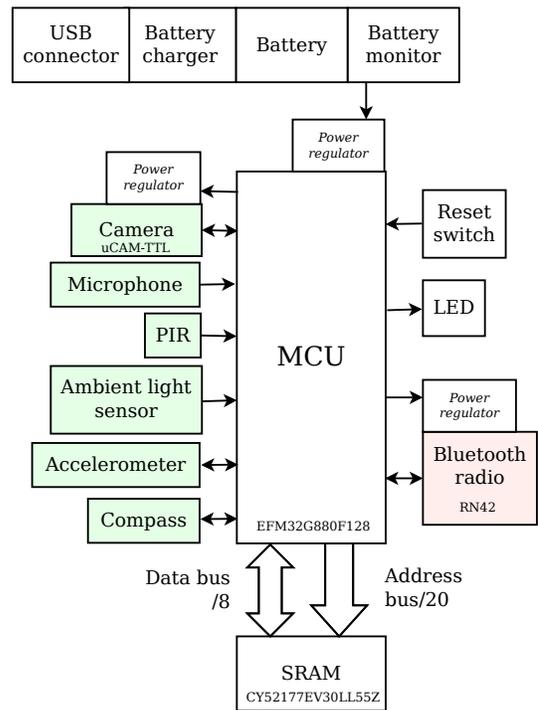


Figure 3: The DejaView device

Table 2: A comparison of sensors considered for the wearable device, with the energy cost per sample and the size of the sensor balanced against potential usefulness. Those in the top part of the table were used in the current prototype.

Sensor	Manufacturer	Device	E/samp.	Vol. (mm^3)	Use
Microphone	Knowles	SPM0208HD5	84.38 nJ	22	Person
3-axis compass	Honeywell	HMC5843	450 nJ	21	Action
Light level	Avago	APDS-9008	672 nJ	1	Place
Accelerometer	Freescale	MMA7455L	39.36 μJ	15	Action
Passive Infrared	Panasonic	AMN32111	24.3 mJ	1621	Person
Heart rate monitor	Polar	RMCM01	10.8 μJ	1416	Health
Humidity	Honeywell	HIH4030	56 μJ	96	Place
Temperature	Maxim	DS1820B	2.25 mJ	36	Place
GPS	RF Solutions	GPS-310FS	742.5 mJ	1142	Place
Blood pressure	Tatung	TMD-36AB	120 J	131040	Health

prototype, some of which are listed in Table 2. Each sensor has been selected to contribute information relevant to the context while adding as little as possible to the power and weight cost. Sensors which can provide information directly relevant to concepts of memory cues (people, places, objects, actions) are valued higher than those which cannot. The energy per sample (as calculated from datasheet values) and the size of the sensor were considered. The humidity, temperature, heart-rate and and blood-pressure signals were judged too energy expensive given that the information obtained did not aid significantly in inferring the context. The sensors used in the prototype are listed in Figure 3. For the microphone, series of samples are filtered and rectified to detect the presence of activity in the vocal frequency range without capturing the actual voice signal. GPS data was added on the phone rather than using an external module, in the interest of power.

To enable advanced rule-based processing of the sensed data, and to add the capability of basic image manipulation, an ultra-low power but computationally capable microcontroller was selected. This, the Energy Micro EFM32G880F128 provides up to 32MHz processing of 32-bit instructions, while only consuming an extremely low $180\mu A/MHz$ (better than most similar 16-bit processors). This microcontroller allows dynamic frequency scaling of the core clock, and can run at several different active energy modes in the system, drawing currents in the order of microamps in most configurations, and hence being very suitable for a low-power design.

As this system relies on a mobile phone, the choice of wireless communications was limited. Mobile phones widely support Bluetooth while other low-power protocols, such as ZigBee, have not yet found the same level

of penetration. Although not the lowest-power technology, Bluetooth does have a reasonably high data rate, which is important for transmitting large photos. A suitable module was selected which featured low-power modes.

The device is designed to be worn close to the user’s eye level, so that photos captured match their view as accurately as possible, helping to provide good cues. This places significant restrictions on the size and weight of the device, in the interests of comfort and unobtrusiveness. As the battery is generally the largest and heaviest component, it is particularly important that the design is conservative with power. Two different lithium polymer batteries were identified and selected for testing with the DejaView device: one of 550mAh, and an additional 250mAh boost battery, for a maximum combined total of 800mAh. A key battery lifetime target is to run a whole waking day, so that the device only needs to be charged overnight. Hence, we aim for a battery life of around 20 hours. The final design weighs 45.6g with the 550mAh battery, or 54.7g for 800mAh. This compares favourably with the 87.6g SenseCam (the current state-of-the-art in wearable memory aids).

Once an image capture is triggered, the sensed information, along with the rules which triggered the capture are encapsulated in an (*information*) *shard* and transmitted to the application running on the mobile phone.

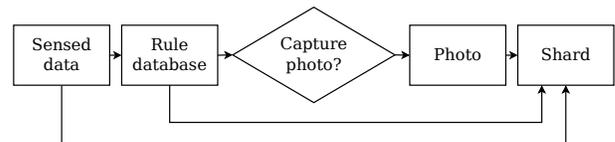


Figure 4: Initial information shard generation process

The appended data is stored in the image file in the Adobe XMP metadata format. This is illustrated in Figure 4.

The application continuously saves photos and sensor data according to the capture rules, as described above. The rules can be adaptively changed through feedback from the phone and Internet service, and can contain (a limited number of) variables, and a small set of pre-defined functions such as sum, magnitude, and time-delayed values. A set lifetime allows the Internet service to tune the device power usage according to the likelihood of capturing interesting images in a certain context.

2.3 Mobile Phone Application

The application running on the mobile phone performs three functions: it processes data from the wearable device according to its own rules, transmits received and additional data (such as GPS) to the Internet service as required, and provides the feedback to the user. The interface component provides the user with a system to rate previously captured information. It also acts as a proxy, connecting the user's online presence to the device and mobile phone application.

The increased computing power of the mobile phone means some more complicated rules can be applied at this stage, such as basic face detection (not recognition); and other simple image processing, such as checking that the lens of the device was not obscured when the photo was taken. Images not judged useful at this stage by these more intelligent classifiers can be discarded, saving power by not uploading unnecessary data to the Internet service.

2.4 Internet Service

The Internet service provides a mechanism for long-term storage, powerful data analysis and a gateway to external online services. The service runs more rules via a data analysis engine, superseding the rules from the previous two levels. At this level, the rules can involve much more intensive processing, including making use of external databases, and external and online services. Information extracted which may prove useful in the future, such as classification of the objects within that environment, is recorded and forms the system knowledge repository. Once the rules have been run, the results of the analysis can be returned to the user, providing the informative feedback. The rule databases on the mobile phone and the wearable device can also be updated. Periodic review is possible via the service at this point.

3. RESULTS

For testing, the device rules were set sensitive to large passive infrared movements, and vocal frequencies detected by the microphone. The mobile phone added GPS data to the information shard, and the Internet service was configured to use the Face.com face recognition service on every photo taken. The names of people recognised were sent back to be displayed on the mobile phone. This was found to function reliably and quickly, validating the three-tier architecture and confirming the practicality of reliable real-time face recognition (Figure 5). Preliminary testing with the IQ Engines image recognition service [17] has also yielded promising results, with the system recognising a range of objects, such as books and crockery, both in isolation and in more naturalistic photos.



Figure 5: Initial system testing, (left) demonstrating autonomous image capture, face recognition, and feedback to the wearer; and (right) the identification of a face in a photo from the device shown the Face.com developer interface

The energy and latency of the system were the key metrics to be evaluated against our original goals of a low-power, real-time aid. The results of this testing are shown below. The time and energy cost could be dramatically reduced by building a custom camera interface and managing the Bluetooth module power more intelligently.

3.1 Energy

In order to calculate the lifetime of the device in different configurations, current draw was measured in different phases of operation. The average time for each operation was measured over repeated runs, and the en-

ergy requirements for the ‘sensing’, ‘capture photo’, and ‘Bluetooth transfer’ operations were calculated using a 4.7Ω high-side current-sensing resistor (Figure 6). In the graph, the sensing, capture, and transfer states can clearly be distinguished, with capture being the most power-consuming, and sensing the least.

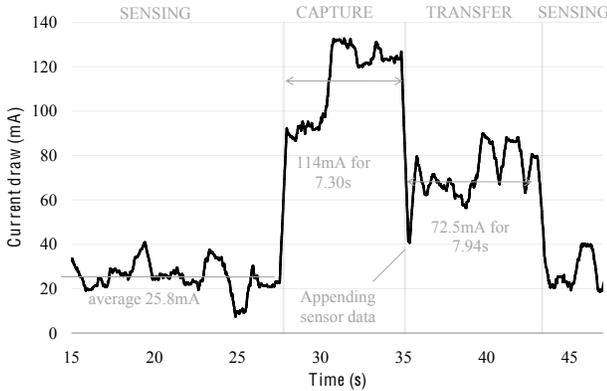


Figure 6: The current draw of the device in various phases of operation

The estimated device lifetime was calculated based on the figures above, and is shown in Table 4. For comparative purposes, we calculated the lifetime with both 550mAh and 800mAh lithium polymer batteries. In most configurations, the prototype device can last a full day on a single charge, and the advantage of a 30-hour over a 25-hour lifetime is minimal as a daily charge will be required in either case. The complete results in Table 3 show the current draw in each phase, the average time taken, and the resultant energy cost.

Table 3: Current draw in the different phases of operation

Phase	I (mA)	t (s)	E (J)
Capture	114	7.30	3.40
Transfer	72.5	7.94	2.36
Sensing	25.8	—	—

While we have shown that the DejaView system represents a step change in the capabilities of wearable memory aids, the most appropriate device for simple comparison of battery life is the SenseCam, as shown in Figure 7. The SenseCam runs slightly longer on a single charge, although it does not transmit images on-the-fly, and so cannot offer interactive assistance. Furthermore, as the DejaView device is currently significantly lighter than the SenseCam, equivalent lifetimes could be reached (with around 1000mAh of battery capacity) while still keeping the device smaller and lighter.

Table 4: Device lifetimes calculated for two battery configurations, for a given average photo capture interval

Average capture interval	Average current (mA)	Lifetime (hours)	
		550mAh	800mAh
15s	53.1	10.4	15.1
30s	39.5	13.9	20.3
60s	32.6	16.9	24.5
120s	29.2	18.8	27.4

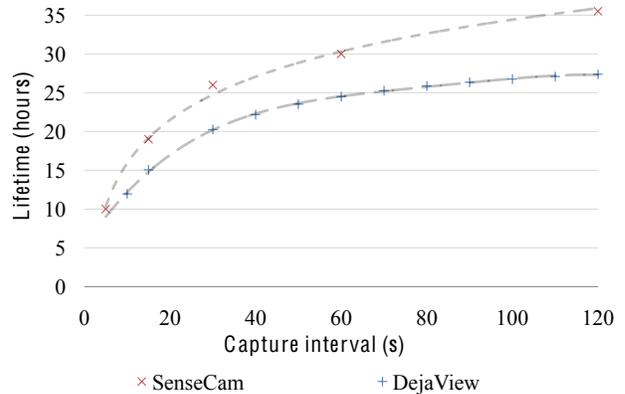


Figure 7: Trade-off between device lifetime and image capture rate for the SenseCam and DejaView device (800mAh battery)

3.2 Timing and Latencies

The latency of the whole system will affect its usefulness: the quicker it is, the more serious the memory problem that can be supported. The most significant latencies in the system were the capture of the photo on the mobile device and the Bluetooth transfer. These accounted for over half of the time taken. The breakdown of the latencies, averaged over a series of tests, is shown in Figure 8. In these tests, the upload from the phone to the Internet was done over a medium-strength 3G connection; improvements to cellular networks and the expansion of public Wi-Fi hotspots will increase throughput, and hence reduce latency.

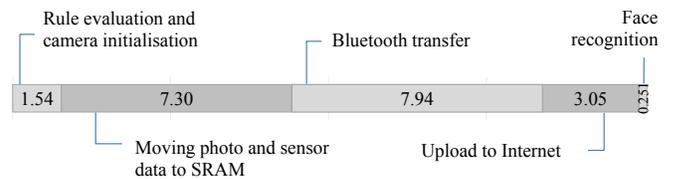


Figure 8: The timings and latencies of each part of the system, in seconds, for a total of 20.1s

At this order of latency, around 20s, the system is useful in a situation where aid is needed for recent episodic memory and details, given more intelligent processing in the Internet service to provide such information. The concept is still being refined, and the existing latencies are currently being improved. As a prime example, the prototype camera module only offers data over a serial interface, leading to a considerable latency in reading photos into memory; once the camera has been replaced, this will be much quicker.

4. CONCLUSIONS AND FUTURE WORK

In this work, we have demonstrated a pervasive, distributed memory aid system which provides prompts in response to the user's environment. The low-power wearable device has been evaluated against the current state-of-the-art and represents a considerable improvement in size, weight, and functionality for a similar power demand. The novel three-tier architecture has allowed for near-real-time face detection to be employed for the first time in a wearable memory aid. Captured photos are available online immediately accompanied by metadata describing the context in which they were taken, for review to support retrospective memory. Future work on the wearable DeJaView device will improve the system latency by several seconds by improving the Bluetooth implementation and the camera interfacing. User feedback will be extended beyond face recognition, implementing the detailed feedback proposed and harnessing ongoing work with novel computer vision systems. Trials are currently being planned for evaluating the system with sufferers of memory loss.

5. REFERENCES

- [1] A. D. Baddeley, M. D. Kopelman, and B. A. Wilson, eds., *The Handbook of Memory Disorders*. John Wiley & Sons Ltd., 2002.
- [2] Center for Disease Control and Prevention, "Public health and aging: Trends in aging," *Morbidity and Mortality Weekly Report*, vol. 52, pp. 101–106, February 2003.
- [3] N. Hersh and L. Treadgold, "Neuropage: The rehabilitation of memory dysfunction by prosthetic memory and cueing," *NeuroRehabilitation*, vol. 4, pp. 187–197, 1994.
- [4] G. Gowans, J. Campbell, N. Alm, R. Dye, A. Astell, and M. Ellis, "Designing a multimedia conversation aid for reminiscence therapy in dementia care environments," in *CHI '04 extended abstracts on human factors in computing systems*, (Vienna, Austria), pp. 825–836, ACM, 2004.
- [5] M. A. McDaniel and G. O. Einstein, eds., *Prospective Memory: An Overview and Synthesis of an Emerging Field*. Sage Publications, 2007.
- [6] M. L. Lee and A. K. Dey, "Providing good memory cues for people with episodic memory impairment," in *Proc. 9th intl. ACM SIGACCESS conference on Computers and accessibility*, 2007.
- [7] M. Lamming and M. Flynn, "Forget-me-not: Intimate computing in support of human memory," in *Intl. Symposium on Next Generation Human Interface*, 1994.
- [8] B. J. Rhodes, "The wearable remembrance agent: a system for augmented memory," *Personal and Ubiquitous Computing*, vol. 1, pp. 218–224, 1997.
- [9] Q. T. Tran, G. Calcaterra, and E. D. Mynatt, "Cook's collage: Deja vu display for a home kitchen," in *HOIT: Home-Oriented Informatics and Telematics 2005*, pp. 15–32, 2005.
- [10] A. Mihailidis, J. N. Boger, T. Craig, and J. Hoey, "The coach prompting system to assist older adults with dementia through handwashing: An efficacy study," *BMC Geriatrics*, vol. 8, no. 1, pp. 28–46, 2008.
- [11] S. Hodges, L. Williams, E. Berry, S. Izadi, J. Srinivasan, A. Butler, G. Smyth, N. Kapur, and K. Wood, "SenseCam: a retrospective memory aid," in *UbiComp 2006: Ubiquitous Computing*, 2006.
- [12] K. Paul-Takacs, C. Moulin, and E. Estlin, "SenseCam as a rehabilitation tool in a child with anterograde amnesia," *Memory*, pp. 1–8, 2010.
- [13] M. L. Lee and A. K. Dey, "Lifelogging memory appliance for people with episodic memory impairment," in *Proc. 10th intl. conf. on ubiquitous computing*, 2008.
- [14] K. Fenwick, M. Massimi, R. Baecker, S. Black, K. Tonon, C. Munteanu, E. Rochon, and D. Ryan, "Cell phone software aiding name recall," in *Proc. 27th int. conf. extended abstracts on human factors in comp. sys.*, 2009.
- [15] M. Mulvenna and C. Nugent, eds., *Supporting People with Dementia Using Pervasive Health Technologies*. Springer-Verlag, 2010.
- [16] N. Caprani, N. Porter, and J. Greaney, "A review of memory aid devices for an ageing population," *PsychNology Journal*, vol. 4, pp. 205–243, 2006.
- [17] "IQ Engines: image recognition and visual search." www.iqengines.com.