MemStone: A Tangible Interface for Controlling Capture and Sharing of Personal Memories

Agon Bexheti, Anton Fedosov, Ivan Elhart, Marc Langheinrich
Faculty of Informatics, Università della Svizzera italiana (USI)
Via G Buffi 13, CH-6904 Lugano, Switzerland
{agon.bexheti, anton.fedosov, ivan.elhart, marc.langheinrich}@usi.ch

ABSTRACT
Today's sensor-rich mobile and wearable devices allow us to seamlessly capture an increasing amount of our daily experiences in digital format. This process can support human memory by producing "memory cues", e.g., an image or a sound that can help trigger our memories of a past event. However, first-person captures such as those coming from wearable cameras are not always ideal for triggering remembrance. One interesting option is thus to combine our own capture streams with those coming from co-located peers, in or even infrastructure sensors (e.g., a surveillance camera) in order to create more powerful memory cues. Given the significant privacy and security concerns of a system that shares personal experience streams with co-located peers, we developed a tangible user interface (TUI) that allows users to *in-situ* control the capture and sharing of their experience streams through a set of five physical gestures. We report on the design of the device, as well as the results of a user study with 20 participants that evaluated its usability and efficiency in the context of a *meeting capture*. Our results show that our TUI outperforms a comparable smartphone application, but also uncovers user concerns regarding the need for additional control devices.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords
Lifelogging; memory augmentation; information sharing; gesture interactions; tangible interface; mixed-method inquiry.

INTRODUCTION
Emerging pervasive computing has made capturing of our mundane experiences straightforward, a practice also known as "lifelogging". Such captured data streams cannot only be used for quantifying our daily routines with the aim of improving our lifestyle (e.g., counting steps, tracking calorie burnout, monitoring sleep patterns, etc.) [4, 29], but can also support human memory [21]. Prior research has investigated how automated capture technologies coupled with data analytics can help us to create "memory cues", a piece of information (e.g., a photograph, a sound, a set of words) that, when reviewed, can trigger recall of previous experiences [13]. Davies et al. define such "pervasive memory augmentation" as a three-step process [10]: captured experience data (step 1) is processed in order to carefully select a set of memory cues (step 2), which are then presented back to users through ambient displays (step 3). Ultimately, the casual review of those cues will allow users to improve their overall recall of the associated experience without the help of any tool.

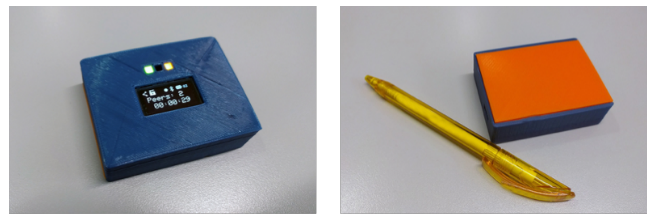

(a) (b)

Figure 1. Overview of the developed MemStone prototype that allows users to in-situ control the capture and sharing of personal memories through a set of five physical gestures. MemStone's front side (a) features a central screen and two LEDs that describe in detail its active operation. MemStone has a different colored back side (b) that allows users to see its position even from distance and denote its action.

In previous work, we built an architecture that not only captures user experiences but also seamlessly and automatically shares them with other co-located users [5]. Our approach primarily focuses on visual logs (photographs and videos), as prior studies identified these to offer the most powerful memory cues [13]. For capturing such logs, we rely on wearable cameras (such as Microsoft SenseCam [19], Narrative Clip, or the recently announced Google Clips), as well as infrastructure cameras. The main driver behind such automatic sharing of memories stems from the low-quality of memory cues captured from body-worn devices [9]: often camera lenses are obscured by hair or clothes; their first-person perspective may fail to capture important details; or they may simply point to the wrong direction. Combining self-captured data with both image streams from infrastructure cameras (their high vantage point allows them to capture comprehensive scenes) and im-

MobileHCI '18, September 3–6, 2018, Barcelona, Spain
© 2018 ACM. ISBN 978-1-4503-5898-9/18/09…$15.00
DOI: https://doi.org/10.1145/3229434.3229477

ages captured from the wearable cameras of other co-located peers (capturing the same experience from different angles) can significantly improve the quality of captured memory cues.

Obviously, a system that enables the seamless sharing of captured experiences introduces significant security and privacy implications [10, 38]. Chief among are the challenges of *(C1) unintended access to captured memories* and *(C2) unintended sharing of sensitive data*. While an experience is captured, the produced memories *should* only be shared with peers that are co-present and engaged with the user. Moreover, while an experience is being captured, there may be moments that users would not want the system to record (e.g., discussing confidential matters in a meeting), or situations that could be captured but which should not be shared with others (e.g., users working in front of their computers). In other situations, a user may be willing to share data with a particular set of users but would want to limit adding further peers. Users should thus be able to control and express their capture and sharing preferences of the event they are experiencing, as the event moves across different levels of sensitivity and privacy. In case that one forgets to react in time, it should be possible for one to perform after-the-fact deletion of such data, as soon as it is noticed. Furthermore, access should be revoked if the 'problematic' data was shared with others.

In our previous work on the architecture for seamless and secure memory sharing [5], we addressed threat C1 above, i.e., preventing passers-by from unintentionally receiving shared experience streams. In this work, our aim is to address threat C2, i.e., to investigate the use of in-situ controls for capturing and sharing captured experiences. To this end we developed *MemStone* (see Figure 1), a prototype of a tangible user interface (TUI) that allows users to control access to (and sharing of) captured memories in-situ. We conducted a user study with 20 participants with the goal of investigating the suitability of a set of gestures to control data capturing and sharing, as well as comparing the usability and efficiency of such interactions with more "traditional" mobile app UI. The study also included an open-ended discussion session to better understand users' perceptions of such tangible interface. This paper describes MemStone's design and functionality, reports on the results of the user study, and concludes with discussing the implications of our results and outlining directions of future research.

RELATED WORK

Our work intersects three principal research strands: privacy issues with visual lifelogs, techniques for enhancing privacy of such data, and gesture-based interactions.

Privacy Issues with Visual Lifelogs

Visual lifelogs can be the basis for highly effective memory cues. Thanks to body-worn cameras, which allow us to seamlessly capture continuous logs of our daily experiences, visual lifelogging lays at the sweet spot between memory recall and ease of capture. However, previous studies have found that such unobtrusive capture can infringe both user and bystander privacy. Clinch et al. [8] have conducted a multi-day experiment where they provided all their participants with a wearable camera which continuously captured users' working activities. Even though they instructed their participants not to record on private spaces (i.e., bathrooms) and placed large "do-not-capture" signs at the entrance of such private spaces, very often users were forgetting to stop capturing in these areas. They also found that such cameras repeatedly captured participants' computer screens and phones. As a consequence, they observed that such capture of private spaces as well as the presence of specific objects in images made users concerned about their privacy. Similar privacy concerns with images showing specific objects or taken at particular locations, but also portraying other known people, bystanders or user activities, have also been observed by other studies [16, 24, 27]. Price et al. [39] noticed that users are less concerned in sharing images with a group of other lifeloggers than with non-lifeloggers, further suggesting that this could re-define what a private space means when lifelogging in a group. In another work, Adams [1] proposes a privacy model that considers *image receiver* and *purpose of usage* as additional factors that exert influence on a sharing decision. All these studies confirm the privacy challenges of visual lifelogging and highlight the need for techniques for privacy-aware data capture and data sharing.

Privacy-enhancing Techniques of Visual Logs

Several solutions have been proposed for regulating access and controlling data sharing [3], however, they usually require active user input in order to specify fine-grained access control and privacy policies. Due to the large volume of captured experience data, this would be a cumbersome process for lifeloggers [8]. For instance, a single Narrative Clip camera produces 120 pictures in one hour, or 1'500 in a day. One option is obviously not to capture all that information in the first place, but the challenge is that one cannot foresee which data might be a valuable memory trigger.

Prior work has made attempts to automate to some extent such decision efforts by designing algorithms that can understand both capture context and captured data. For example, Fan et al. [14] propose a mobile-based technique that stops lifelogging capture when it detects that a user is in a restroom. Moncrieff et al. [34] leverages background audio, but also other sensors, to determine the context in surveillance scenarios running in private environments, such as smart homes. Based on the inferred context, the system will activate a predefined privacy policy and enforce it using a combination of data hiding techniques. Other approaches rely on computer vision algorithms to study the captured images themselves and flag those that contain specific places [43], particular objects [15], computer screens [28], or even images that portray activities [31].

While all these (semi-)automated solutions can potentially improve user privacy, however, they may reduce the utility of a memory augmentation system. As Adams [1] notes, privacy concerns related to captured experiences often rely on users' implicit assumptions of its usage and intended receiver, and as such they can vary with person and context [8]. For instance, an image that can infringe a user's privacy because it contains a computer screen can be *the* strongest memory cue; or a user might want to share such computer screen image with only a particular other user that she trusts more. In contrast, we propose a solution based on (manual) in-situ user input. In-situ

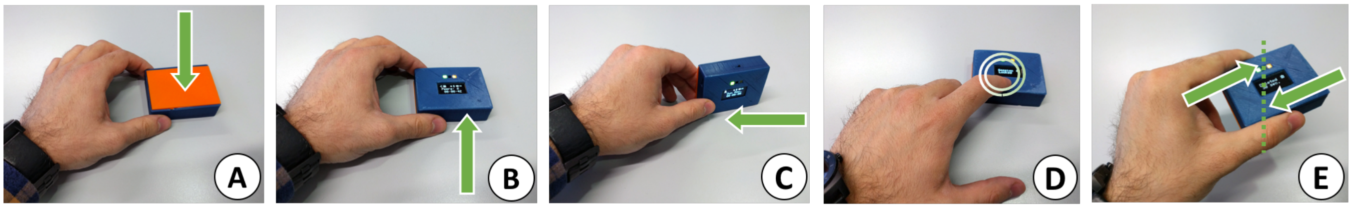

Figure 2. The current MemStone prototype supports five physical gestures: (1) *face-down*; (2) *face-up*; (3) *stand-on-side*; (4) *double-tap*; (5) *shake*.

controls offer greater flexibility to users and allows them to react in real-time based on their impressions of the context, but still keep user involvement lower than post-hoc solutions. Hoyle et al. [23] also confirm that lifeloggers prefer in-situ control more than manual post-hoc filtering. Ultimately, our in-situ control can go side-by-side with an automatic control approach and complement it.

Gesture-based Control Interfaces

Prior research has explored the opportunities of controlling virtual information using objects from the physical world. Fitzmaurice et al. [18] propose a technique for manipulating digital data using graspable wooden blocks, with the goal of augmenting traditional graphical user interfaces. The *tangible bits* vision by Ishii and Ullmer [25] aimed at bridging the gap between digital bits and graspable objects, where objects from the physical world would both manipulate and visualize digital content. Similarly, Fishkin et al. [17] present the paradigm of embodying physical manipulations to computational devices, so that the device's physical body becomes also its interface.

Other research has particularly focused on box-shaped physical interfaces. For example, Rekimoto and Sciammarella proposed ToolStone [40], a cordless tangible interface controllable by physical manipulations. ToolStone would be operated by users' non-dominant hand and would complement the traditional computer input device (i.e., the mouse) in various applications where such a bimanual interface could be appropriate, e.g., choosing a color from a palette; zooming, scrolling or rotating screen contents; controlling a virtual camera, etc. Sheridan et al. [41] explored the *affordance* of a cube as control interface. Through a user study they developed a classification of 16 distinct gestures (or non-verbal dynamics) that users performed with a cube, such as placing the cube in a particular place or position, turning it, rotating, tapping, shaking, squeezing or fiddling with it, etc. Van Laerhoven et al. [45] build a such a cube that embodies gesture recognition and show how it can be used as an input device for desktop applications involving selection and navigation operations.

All these studies show the feasibility and psychological affordance of box-like interfaces, however they mostly focus on applications for extending traditional input devices or GUIs. Moreover, in these works the box was used only as an input device and was not utilized to also provide feedback back to the user. In our work we adopt the concept of a box-shaped interface and apply it in a scenario that goes beyond extending conventional input devices, i.e., allowing one to control and observe *how*, *when* and with *whom* one's lifelogging devices are capturing and sharing data that constitutes one's memories.

MEMSTONE INTERFACE

We designed MemStone inspired by Rekimoto's and Sciammarella's ToolStone [40], particularly by its design and shape, but also based on design requirements by prior work on tangible interactions [25, 44]. MemStone is a rectangular-shaped 3D printed box (measuring 67 $mm \times 52$ mm); augmented with an embedded computing platform (NodeMCU ESP8266[1]), an accelerometer sensor, a vibration mini motor, wireless communication capabilities and a lithium battery. In principle it can thus communicate directly with a range of capture devices (e.g., body-worn cameras, smartglasses, audio recorder bands, etc.) in order to allow users to in-situ control, with simple physical manipulations, what memories such devices can capture and share with co-located others.[2]

Its shape (rectangular) and visual appearance (bi-colored) aim to allow users to easily understand its current operation (to some extent also its available gestures) at a glance. The front side (see Figure 1–a) has a central screen (with a resolution of 128×64 pixels and a diagonal size of 33 mm) that provides feedback on the system's current mode of operation. Specifically, the screen shows several aspects of the active action, such as elapsed time of current data capture (if such practice is active), the number of peers one is sharing data with (if any), if newly joining peers are allowed to automatically get a copy of captured data, and the device's remaining battery level. The front side has two additional LEDs that also give details about the current operation. The LED on the left will signal the user when the experience is being captured while the other LED will indicate that such data is also being shared. The back side (see Figure 1–b) has a different color than the other sides to allow the user, as well as other co-located people, to see the device's state from a distance and thus allow them to note its active operation without having to closely look at its screen.

Gestures and Control Actions

Starting from the previously described challenges, we have derived five different aspects that one could control when capturing and sharing experience data. Each such control action can be executed by performing a particular physical gesture using MemStone, as shown in Figure 2. For the gesture selection we were in part inspired by the work from Sheridan and her colleagues [41]. Through a user study, Sheridan et al. explored the natural affordance of a cube-shaped device and came up with a classification of 16 physical gestures that

[1] http://www.nodemcu.com

[2] Note that, however, the current prototype of the MemStone features gesture recognition but it does not implement a protocol for communicating with other devices.

a cube affords. Starting from their result, we selected 5 such physical gestures that we believed to offer a good match to the actions for controlling experience data capture and sharing:

- The ***face-down*** gesture (Figure 2-A) stops both *data capture*, sourced from any of the user's capture devices, and *data sharing* with other co-located people. This way a user can let others know that she does not record anything herself. While this may also signal that a user does not want others to record her, the MemStone cannot control the capture operations of other users.
- By putting the MemStone device ***face-up*** (Figure 2-B), the user triggers her data capture from any of her lifelogging gear. In addition, the user informs co-located peers that she is recording and expresses her willingness to exchange data with them. Sharing commences automatically with all other peers that have similarly positioned their device face-up.
- One could as well capture data from their own recording gadgets but inform others that she does not want to share data of that particular moment. This is achieved by putting the MemStone on a vertical ***stand-on-side*** position, facing oneself (Figure 2-C). Any active sharing session with any peer stops immediately.
- ***Double-tapping*** the MemoryStone (Figure 2-D) "locks" the data exchange with the current set of co-located people and prevents any further peers from joining (though peers leaving will still be removed from the common data exchange). A subsequent double-tap will remove this lock. A similar access control mechanism, but using the metaphor of a virtual wall, has been proposed by Kapadia et al. [26].
- The ***shake*** gesture (Figure 2-E) allows the user to delete the last 30 seconds (this is configurable) of captured data. By repeating this gesture the user can delete data captured for longer periods.

Gesture recognition is based solely on data obtained from a 3D accelerometer sensor. Triggered actions are confirmed with the help of distinct vibration patterns as feedback. For the lock/unlock action, two different vibration patterns will signal the user the current lock state after the initiated change.

Envisioned Usage Scenario

To better illustrate the vision of (i) a pervasive memory augmentation system that also supports seamless sharing of memories and (ii) the MemStone's concept as a tangible control interface for such a system, consider the following *meeting capture* scenario, as shown in Figure 3.

Team Alpha, with its members Bob (Figure 3 left), John (Figure 3 middle) and Alice (Figure 3 right) have recently decided to use a memory augmentation system during their weekly meetings. During one meeting, Bob and John use a body-worn camera that automatically captures an image every 30 seconds. Alice uses a set of smartglasses that, similarly to the cameras, capture 2 photos per minute. Bob also brought a wristband that captures the last 30 seconds of audio with a single tap.

Each device uploads the data it captures to the respective user's memory repository, where such data is then processed in order to generate memory cues. Eventually, generated cues will be

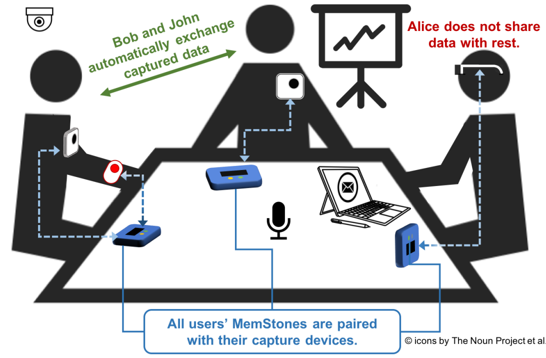

Figure 3. Illustartion of the envisioned usage scenario.

delivered to the user through ambient fashion displays (e.g., screensavers, phone lock screens, picture frames, etc.). By reviewing such cues, they would improve their recall of past meetings and thus be more prepared for their next meeting.

All team members use the MemStone device to control their meeting capture and sharing practices. At the scenario in Figure 3, Bob and John have positioned their MemStones *face-up* notifying others that they are capturing this part of the meeting, as well as expressing their willingness to exchange data with the rest. As a consequence, they would get access to images captured from each other's cameras. John is particularly interested in obtaining images from Bob's camera, especially since his camera could occasionally capture the whiteboard (which is not the case with John's camera). In addition, John would also obtain audio snippets recorded by Bob's wristband. Alice, on the other hand, fears that her glasses might capture some sensitive information from her laptop's screen, which she has in front of her. She decides not to share any data with the rest, but only records for herself. To do that and also to inform others that she is not willing to share anything, Alice has put her MemStone on the *stand-on-side* position.

In this scenario, we envision that MemStone devices will be part of the meeting room, just as one finds a remote controller for the projector or whiteboard markers in such a room. Before the meeting start, each attendee picks up a MemStone from a small receptacle at the entrence and connects it with their meeting capture gear (e.g., by simply physically touching those devices together [22]). The connection would remain active as long as the devices are in close proximity. After the meeting ends, users would return the MemStones to their receptacle or simply leave them on the table. Once a user leaves the room, the MemStone disconnects from their capture devices and reverts back to an idle state.

Using MemStone to Control Infrastructure Sensors

In addition to data sourced from their personal devices, team Alpha members occasionally use the room's built-in capture infrastructure to also make recordings of their meetings. The room is equipped with a fixed-camera, a central audio recorder, as well as a smartboard that captures snapshots of its contents.

Meeting attendees could use their MemStone to implicitly control experience capture from such infrastructure devices. A strict privacy-aware approach would stop recording from any fixed sensor as long as there is at least one participant that has positioned their MemStone face-down (i.e., does not record for themselves and does not want to exchange any data with others, Figure 2-A) or also on a vertical stand-on-side position (records only for themselves without sharing any data with others, Figure 2-C). Another approach would be that recording is based on a majority decision, or even unanimosity (i.e., infrastructure sensors stop recording only if *all* users disable recording). Alternatively, the room's capture devices could also be controlled by a designated "room-MemStone". Its operation would then need to be agreed on by all participants.

In this work, we consider a similar scenario (i.e., a meeting capture between two attendees in which data is sourced only from user controlled devices) in order to: 1) investigate how controlling data capture and sharing practices using physical gestures compares with a more traditional alternative (i.e., smartphone app), and 2) understand user perceptions regarding the MemStone gesture-based interface.

PHONE-APP AS AN ALTERNATIVE INTERFACE

Prior research in lifelog privacy [23] suggests that lifeloggers prefer in-situ control over end-of-day filtering for deciding which visual logs are shareable with others. Traditionally, such in-situ privacy control approaches are implemented as mobile apps running on a touch-based smartphone [2, 7]. While emerging body-worn devices, such as smartwatches or wearable glasses, could in principle become viable alternative platforms for designing privacy control apps, the widespread use of phone apps means that a smartphone-based tool represents the most viable alternative to our proposed tangible interface today.

We thus created a simple smartpone app that can act as a baseline alternative to MemStone. The phone UI, displayed in Figure 4, allows one to perform the same five control actions that can be performed using a MemStone. Note that since we wanted to compare the gesture-driven MemStone against the idea of using the smartphone as control interface (and not compare with a particular interface design per se), we went with a rather simplistic but intuitive UI: operated by buttons, each with a short label clearly describing its action. In the previous meeting capture scenario, users would install this app on their personal phones and pair it with their lifelogging gear.

USER STUDY

We conducted a user study aimed to address the following research questions:

1. Is the proposed gesture-based interface usable for in-situ controlling data capture and data sharing in meeting capture scenarios?
2. Are the chosen in-situ controls easily remembered even after longer time periods?
3. How does the use of gesture-interactions to control data capture and sharing perform against conventional (i.e., mobile app-based) user interfaces?
4. What are user perceptions on such gesture-based control interface within the context of memory augmentation systems?

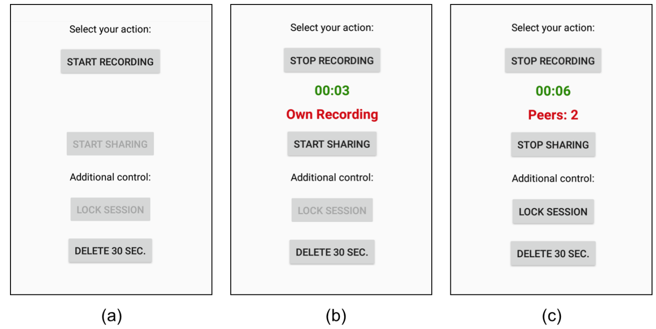

Figure 4. A phone-based interface for controlling data capture and sharing of daily experiences. Screenshots of different interface states: a) initial state when no capturing or sharing is taking place; b) state of an ongoing data capture but not sharing; c) interface's state after activating data sharing with co-located peers.

For this study we recruited 20 participants (7 of them were female) using snowball sampling. Their age ranged from 22 to 63 years old, with an average age of 28.75 years (SD = 9.31). They had different education levels, 4 of them had a high-school degree, 9 were bachelor graduates and 7 had a masters degree. Most of our participants stated that they have an affinity for technology. All said that they used a smartphone several times a day and a laptop few times a week. No incentive remuneration was provided to study participants.

Study Design and Procedure

We performed a comparative user study, employing a within-subjects design in a counterbalanced order, where each participant tried both interfaces, MemStone and smartphone app, in differing order (i.e., MemStone–Phone or Phone–MemStone).

We considered a meeting capture scenario between two people, similar to the previous scenario in Figure 3. To strike a balance between validity and repeatability, we prepared two different videos, each showing a meeting. We then asked participants to watch those videos and pretend to be one of the attendees. Each such recorded meeting contained 10 tasks related to controlling capturing and sharing of memories within the meeting (note that a detailed description of both meeting videos and tasks is provided in the next section). Prior to watching a meeting video, participants were given either our MemStone or an Android mobile phone with the corresponding control app installed (accessible via a shortcut from the home screen). They were then instructed to use the respective tool for handling these control and sharing tasks. The phone was a Nexus 5X running stock Android 8.0 with only our app additionally installed. Note that the Nexus 5X features a default screen timeout of 15 seconds, which we did not change. While no screen lock (e.g., PIN) was setup, participants had to turn on the phone and then swipe up the lockscreen before they could interact again with the app after a screen timeout.

A session with a single participant lasted on average 60-70 minutes. At the outset, we briefly introduced the study, asked participants to sign a consent form, and to provide basic demographic information. Then, the session proceeded with the following stages:

1. Introduction of the vision for technology-driven memory augmentation, with a focus on how it could be applied to the envisioned meeting scenario.
2. A short demonstration of one of the interfaces followed by a trial session where participants try the different functionalities of the chosen interface.
3. Participants watch a video of a recorded meeting and use the assigned interface to control access to (and sharing of) their meeting memories.
4. Users fill a SUS questionnaire to express their perceived usability of the interface.
5. Repeat steps 2–4 using the other control interface and the second video.
6. Semi-structured exit interview, reflecting on the experience with both interfaces.

Participants' interactions with the interfaces (steps 2–4) were video recorded using a wide-angle camera. The produced video data was later used to compute the devices' efficiency and effectiveness for performing the specified control actions. The goal of the semi-structured interviews was to better understand the user experience with both control interfaces, and to also explore user perceptions on the proposed in-situ control interface. We recorded these sessions using a voice recorder, and then transcribed the recorded interviews. To analyze this data, we followed an iterative process, going back and forth between the data and the researchers' notes [33]. This technique helped us to organize participants' feedback related to our TUI's physical design, its interaction, as well as participants' perceived usefulness of the system.

Recorded Meetings and Tasks

To better simulate a real meeting scenario, we created two videos depicting a meeting between an instructor and a teaching assistant, in which they discuss the progress of a course they teach together. In the video, both attendees capture the event using a wearable camera, while one of them also uses a wristband audio recorder. During the video, the attendees discuss both non-sensitive and sensitive issues (e.g., student grades), thus requiring several control actions on the capture and sharing system. Each time such a control point comes up, the "actors" explicitly announce this need for control (e.g., "Let me pause recording for a moment!" or "Can you delete this part, please?"), yet without explicitly stating what exactly they have to do. A small icon at the bottom-right corner of the screen additionally triggers the need for action. To allow for different participant reaction times, the study administrator would then remotely pause the video and resume it only after the participant would perform an action, or would say that they would not know what action to perform. For each action, we measured participants' reaction time and task completion rates.

In total, there were 10 tasks involving all five control actions that were presented previously (3 of them were repeated twice), plus two additional information gathering tasks related to the feedback given by the control device (see Figures 2 and 4):

T1: *Capture and share with co-located peer*s. MemStone: face-up; Phone: press "Start Recording" and then "Start Sharing" buttons.

T2: *Check elapsed time of current capture.* MemStone: shown on small display. Phone: shown in app.

T3: *Stop capturing and sharing*. MemStone: face-down; Phone: press "Stop Recording"

T4: *Capture only for oneself and do not share with others.* MemStone: stand-on-side; Phone: press "Stop Sharing" if currently sharing.

T5: *Capture and share with co-located peers* (see T1).

T6: *Delete the last 30 seconds of captured data.* MemStone: shake; Phone: press "Delete 30 Sec.".

T7: *Lock current sharing session.* MemStone: double-tap; Phone: press "Lock Session".

T8: *Verify with how many peers the system is sharing data.* MemStone: shown on small display; Phone: shown in app.

T9: *Stop capturing and sharing* (see T3).

T10: *Capture only for oneself and do not share with others* (see T4).

The produced meeting videos are slightly different in order to minimize participants' learning bias when they have to watch both videos for trying the two interfaces. In addition, the videos also feature different task sequences.

RESULTS: INTERFACE COMPARISON

Efficiency and Effectiveness

Initially we compared both interfaces in terms of *effectiveness* (i.e., task completion rate) and *efficiency* (i.e., task completion time). For each device we collected data from 10 interaction tasks from a single participant, resulting in a total of 400 tasks performed from all 20 participants with both control interfaces.

For computing the task error rate, we only looked at the human error aspect and did not consider mistakes from the system (e.g., if the MemStone device failed to recognize the performed gesture). As displayed in Figure 5–a, participants on average performed equally well with both interfaces. However, when looking at individual tasks, MemStone users performed less mistakes than phone users in tasks related to 'capturing and sharing with others' (T1 and T5) and 'capture only for oneself' (T4 and T10). The challenge with performing tasks T1 and T5 using the phone UI is that one has to press two buttons, one for recording and for sharing, as opposed to the single gesture (face-up) with MemStone. From our observations, we saw that most participants pressed only the "start recording" button in these cases. In two cases, we observed that participants confused the MemStone gesture for performing T4 and T10 (i.e., stand-on-side) with those of T3 and T7 (i.e., face-down and double-tap, respectively). When using MemStone, two participants failed to do T2 (reading the elapsed time of the current capturing) and T8 (verifying that that the preceding task on locking the sharing session) since they did not notice that such information was displayed in the MemStone's screen.

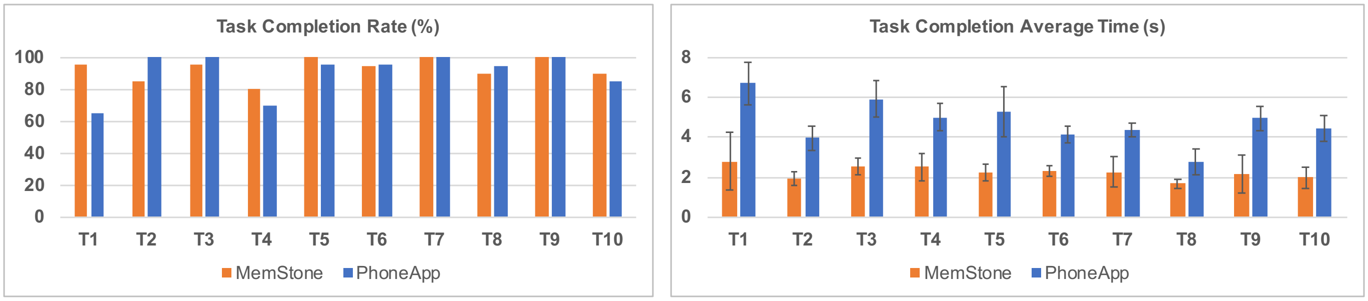

(a) (b)

Figure 5. Effectiveness and efficiency results: a) task completion rate in percentage; b) average task completion time in seconds (includes data from successfully completed tasks only). Error bars represent 95% confidence interval.

For each task we also measured the time between the moment the visual clue was provided (an icon being displayed in the video of the recorded meeting) to the moment participants performed the correct action. Such information was precisely computed from the session video recordings using the timestamps overlaid on those videos. For all tasks, MemStone outperformed the phone interface in this aspect: using the MemStone participants could compute a given task in average 2.5 times faster than when using the phone interface (see Figure 5–b). We conducted a repeated measures ANOVA to compare the effect of control interface (MemStone and Phone app) on task completion times. Results show that MemStone was significantly faster than Phone app (Wilks' Lambda $= 0.485$, $F(1,180) = 190.799$, $p < 0.001$).[3] Moreover, we performed a Sobel test [42] to check if these results are not due to the fact that some of our participants had not recently used an Android phone (and that they could have spent more time performing the study tasks using the Android-based test phone, thus, making the control with the phone slower). Test results confirmed that the effect of control interface on task execution times was not significantly mediated by participants' smartphone ownership ($z = 0.056$, $p = 0.954$).

This was also reflected in participants' perceptions during the interview sessions, where the MemStone was perceived to be more efficient in performing the given actions: *"I think the dedicated device is better, it is just easier. You do not need to check your phone and you are not losing time. While in the phone maybe you get a message and you want to read it, hence it is not efficient."* (P13, similarly P14, P19, P20).

Perceived Usability and Learnability

We also compared the two interfaces with regard to users' perceived usability and learnability. After watching a recorded meeting and using one of the two interfaces for controlling meeting capture and sharing, participants evaluated the interface using a SUS questionnaire. SUS is a 10-item questionnaire that has been extensively used by usability practitioners for assessing only the perceived usability of a system. However, two recent works [6, 30] suggest that the SUS result can be decomposed into two components for measuring both usability and learnability of a system. After applying this approach, the average usability scores for both the MemStone and the phone interfaces are 75.31 (SD=18.38) and 82.81 (SD=16.42), respectively. This suggests that the gesture-based MemStone prototype is in principle usable above average. However, it echoes user concerns regarding the necessity for an additional control device, which we discuss later in this paper.

As for the learnability aspect, MemStone scored an average of 78.75 (SD=22.61) and the phone's score is 89.37 (SD=13). Unsurprisingly, the phone – being an already established concept and an artifact that most of us know how to use and operate –, scored higher than the novel interface concept of MemStone. All phone buttons were also unambiguously labeled, making it easy to use without having to remember much. However, the small difference in their scores implies that participants find MemStone not much harder to learn. This was also underlined from some participants during the interview. They believed that MemStone can be quickly learned also by children: *"I think it was a bit easier to use than the phone app. You do not have to unlock your phone and choose the right button, so it was quite intuitive and user friendly. And I think even children would be able to use this correctly and learn it in few minutes."* (P2, similarly P7, P14).

Perceived Intuitiveness and Enjoyment

During the interview sessions, we asked participants to reflect on how they would compare these interfaces in terms of intuitiveness and enjoyment. Enjoyment is an important part of the usability of a product, as it positively influences both a user's willingness to learn and their tolerance for interface shortcomings [37]. Participants acknowledged the fact that the phone UI was clear and intuitive, and that it was similar to many other apps that they use everyday: *"The phone is more intuitive. You have the feedback you know exactly what is going on. As for the buttons you know what each of them does."* (P13, similarly P19). However, they also regarded the MemStone to be equally intuitive to interact with: *"The phone is labeled, something you use everyday. It is easy and more intuitive. It does not mean that the box was harder, it was also really easy to understand."* (P19, similarly P14, P16, P20). Some others would even believe that MemStone can actually become more intuitive once users get to know it better: *"Once you pass the learning phase it becomes more mechanic, probably without thinking so much you would just use the MemStone."* (P17, similarly P5).

[3] Statistics were computed using data only from successfully completed tasks.

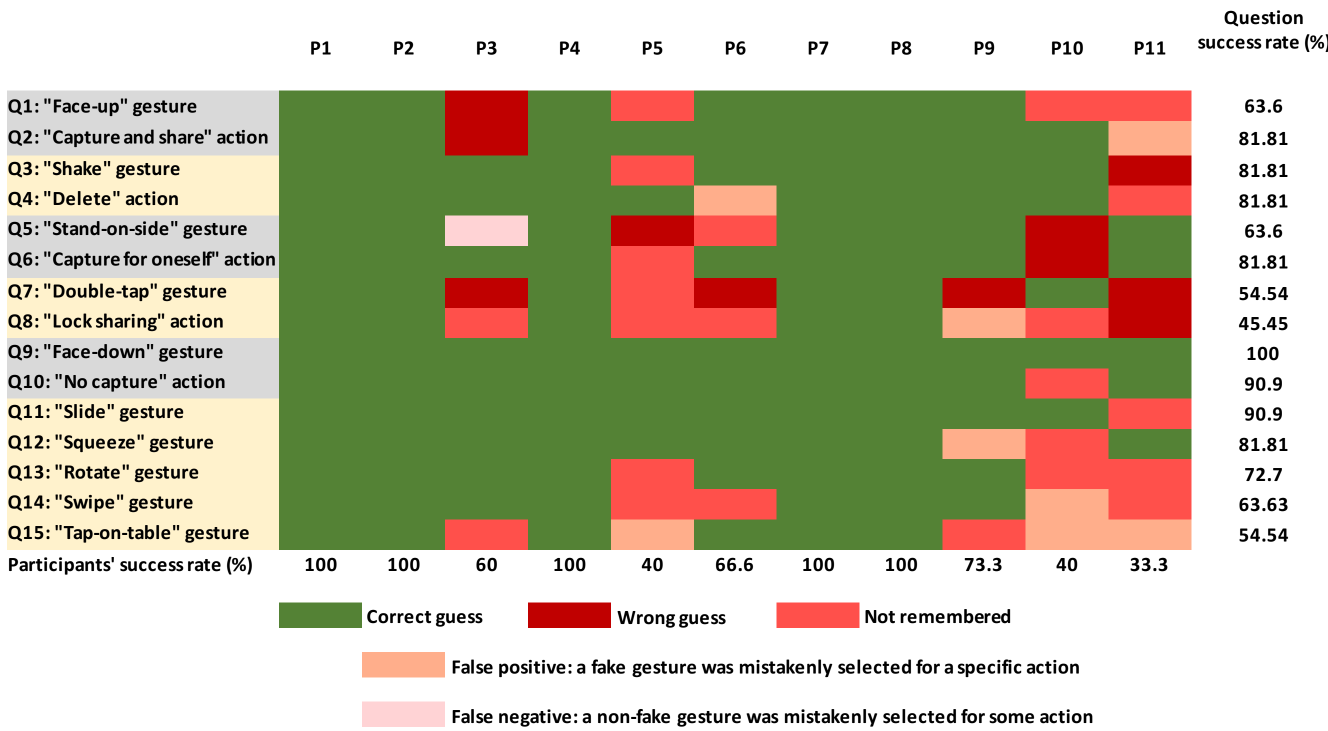

Figure 6. A heatmap illustration of the distribution of answers from the survey on long-term gesture memorability. Odd numbered questions from 1 to 10 ask participants to map a given gesture to an action, while even numbered questions ask the opposite. Questions from 11 to 12 quiz participants on the action that would be triggered from a fake-gesture that was not used in the first study.

Participants were less divided in expressing their opinion on devices' perceived enjoyment. A clear majority said that using the MemStone was more enjoying and fun (P1 but also P2, P4, P6, P9, P13-P20). The reasons for such choice were various: the phone UI was *just* another app (P13, P14, P16); the MemStone was a novel concept operated by physical manipulations and is something that one will not encounter frequently (P13-P20); or also because the phone is more invasive and requires active user focus: *"The phone took out the fun the moment I had to stop whatever I was doing and focus on it to search the necessary functions. That was not fun! The cube was definitively more fun."* (P20, similarly P17). Other participants compared the experience of using the MemStone with that of a toy: *"The MemStone is more fun to use, it is a toy essentially. . . "* (P13) while P20 got emotionally attached to it comparing it with a pet: *"And I like the vibration, it becomes like pet in a way. I guess you can get very emotionally bonded to it."*.

GESTURE LONG-TERM MEMORABILITY

In addition to exploring several characteristics of MemStone gestures, including efficiency, enjoyment, learnability and usability, we also explored gesture *memorability*. Prior research highlights memorability as a key characteristic of gesture-based interfaces [20, 32], since easy to remember gestures can reduce mistakes and frustration (especially when one is focused on other important things during a work meeting). Additionally, memorability can also increase adoption of a gesture-based interface [35].

After conducting the comparative user study, we administered a follow-up study with the goal of investigating participants' long-term memorability of the MemStone's physical gestures. This follow-up inquiry was conducted four months after the first study. We contacted all prior participants by email and asked them to participate in a short follow-up online survey. 11 participants from the prior study (55%) participated in the follow-up study. As in the first study, no incentive payment was provided to participants for this follow-up study.

Follow-up Study

The survey contained two types of multiple-choice questions. The first 10 questions showed a MemStone gesture (e.g., "face down") and then asked participants to select which of the given five actions would be triggered after performing the gesture. For each gesture the survey showed a small embedded video of how it is performed. As the original study only had five gestures for five actions, we included five *fake*-gestures that were not used in the previous study, yet which were somehow similar to the original five gestures: squeeze, rotate, swipe, tap-on-table, and slide-device-back-and-forth. The fake gestures were meant to understand how memorable the original gestures were. Given that not all of the 10 gestures mapped to an actual action, participants could answer "This gesture was not used in the study", as well as "I do not remember this gesture". A second part of the survey questions then investigated the reverse, i.e., the mapping of actions to gestures. Here, 5 questions showed an action (e.g., "pause recording") and participants then had to select which of the 10 previously shown gestures (i.e., 5 real gestures and 5 fake gestures) would trigger this action. An 11th option again was "I do not remember the gesture for this action". All 15 questions were presented to participants in random order.

Results

Figure 6 provides an overview of the results from the long-term gesture memorability survey. From participants' individual responses (columns P1 to P11 in Figure 6) we observe that five participants answered all questions correctly, three participants have a correct response rate between 60% and 73%, while the other three participants (P5, P10 and P11) could only successfully answer less than 60% of the questions. We investigated further their individual performances achieved during the first study. There was no indication for P5's low performance, however, for both P10 and P11, their low level of engagement during the first study and also their lack of interest for extra electronic devices seem reasonable explanations for their low performance. This can also be confirmed by their perceived SUS scores for MemStone's usability (32 and 55).

When looking at the gesture-action mapping (see rows Q1 to Q10 from Figure 6 with a combined order of gesture and related action questions), both the 'face-down' gesture and the 'no capture' action were successfully mapped in 90% of the cases. The second best such mapping (with a correct response rate of 72%) is for the 'shake' gesture and 'delete 30 seconds of data' action. The weakest mapping with only 45% was for the 'double-tap' gesture and the 'lock sharing' action. By further investigating this association, we observed that in four cases the 'double-tap' action was mistaken with another action (once with 'delete 30 seconds of data' and three times with 'capture and share data'). When asked to identify the action for the 'double-tap' gesture, four participants said that they do not remember it, and two others selected a wrong gesture ('face-down' and the 'squeeze' fake gesture).

Among the fake gestures, i.e., physical gestures that were not used in the first study (rows Q11 to Q15 from Figure 6), the 'tap-on-table' and the 'swipe' were wrongly selected as to trigger an action in 3 and 2 cases, respectively. In line with the previous outcome, the 'lock sharing' was the mostly mistaken action to be triggered by such fake gestures.

All in all, even after four months participants were able to successfully recall the relationship between most physical gestures and control actions. Apart from the association of the 'double-tap' gesture and 'lock sharing' action, all others gesture-actions were correctly identified in more than 60% of the survey questions, which suggests that they can be successfully used with the proposed tangible interface concept.

USER PERCEPTIONS OF THE MEMSTONE DEVICE

Suggested Design Improvements

Participants generally liked the concept of the rectangular-shaped box. They also liked its bi-colored design, the two LEDs (which indicate capturing and sharing activities), and its central screen. However, they proposed several improvements related to its physical design.

Participants believe that making it *smaller* and *lighter* would also make the device to better fit their hands, thus being also easier to be used (P2, P13). Moreover, it would also become less of a burden to carry it with them: *"I think you can do this device smaller, like a USB stick. I have 3-4 sticks with me and it is not a big problem."* (P13, similarly P14). Additionally, the current prototype was perceived as being a bit slippery, so designers should take into account materials which provide better grip: *"The current prototype is a bit slippery and it can easily escape from the hands."* (P15).

To further improve the device's feedback, participants suggested to include an additional LED for the 'lock sharing' action, similarly as for the 'capturing' and 'sharing' actions. Participants expressed contrasting opinions regarding the amount of information displayed on the MemStone screen. Some would like to have a less cluttered screen that would show only the number of peers one is sharing data with, together with the elapsed time of sharing, while others would want even more information, such as as detailed profile information on the connected peers and their physical distance. This suggests that one should consider a customizable interface that can be switched between such simple and comprehensive views.

Interaction Techniques

Gesture Affinity

Study participants expressed an affinity for the underlying physical gestures. The mapping between gestures and their control actions was well aligned to users' mental models, thus it was easy to associate them during the study: *"I think it is very intuitive, if you have it face-up, it kind of radiates through versus everybody sitting around, facing yourself it is just you and face-down is off. You have chosen the functions nicely with the physical movements and the device's position."* (P7); *"Once you rationalize them, you see that some are quite easy, such as turning on the other side. But also other ones pair to their actions and make absolutely sense. From that point of view they become easy to remember and can be used without thinking."* (P16, similarly P9, P14, P17-P20). P16 further commented on the low physical demands for performing such gestures, suggesting that the interface can be also operated by people with slight physical disabilities.

Gesture Challenges

Users also expressed some concerns and challenges with the 'shake' gesture, as well as with the 'double-tap' one. While the mapping between 'shaking' and 'delete data' was not necessarily questioned, it was seen as a rather hard-to-perform gesture (P4, but also P5, P7, P14, P17, P18). Since most of us tend to also move our hands while we speak, it is likely that the system can mistakenly delete our data in such cases. Moreover, it was suggested that it may be even considered as not polite to perform it while one is speaking in a meeting. Suggested alternatives were 'swipe' or 'squeeze'.

While 'double-tap' was considered an easy-to-perform gesture, most participants expressed their concerns regarding its relationship with the 'locking sharing' action. They also believe that 'locking' should be extended beyond a binary lock/unlock model, so that one can be more selective on whom to keep and remove from their locked session.

Role of Device Visibility

Participants appreciated the fact that MemStone is more visible and transparent than the phone in conveying its feedback to all co-located users. Such openness had a two-fold effect on participants' perceptions about the device. First, participants said that they felt more confident about their privacy. For instance, should users agree not to record some part of the meeting, then by looking at others' MemStones they can easily understand if they are behaving according to the agreed protocol: *"If there's something that doesn't have to be recorded and I ask for it, then I can see if people are following. This gives me more confidence. It's a way to see what people do and how they behave; it's an additional feedback."* (P15, similarly P4).

The increased confidence, however, may also come with a cost. Some participants said that the device's transparency could also influence their data capturing and sharing practices: *"It would extremely change my behavior even if I might try not to let it influence me. Thinking that somebody is watching me, or even recording me, you become self-aware, it changes your behavior, you want it or not."* (P17, similarly P2, P4, P14). However, others suggested that their decision to share or not actually depends on the context and on the other attendees (P13, P15, P18, P19). This outcome is in line with prior work [46] which suggests that knowledge sharing behavior is influenced by multiple factors. Lastly, some participants expressed their belief that by observing the action of others one can increase their meeting concentration: *"Since one is recording, then maybe someone will say something important and I should record too.'* (P14, similarly P16, P17).

DISCUSSION

Overall, our TUI was efficient in allowing users to control access to, and sharing of, captured experiences in a meeting context. It was also perceived as user-friendly and enjoyable to use, which is in line with findings from Hoyle and colleagues [24], suggesting that users preferred an in-situ control method rather than other post-hoc approaches, such as [15, 43].

MemStone's efficiency and effectiveness, but also our participants' affinity to its gestures, highlight that the device's affordances (i.e., range of possible activities) are visible and clear to users. This is in line with findings from Sheridan and colleague [41] on the affordances of box-shaped interfaces, also following Norman's insight that such affordanes are useless if they are not visible to users [36]. We also found that participants could relatively well remember the physical gestures and their corresponding actions even when asked four months after. Our results suggest that a reasonable relation between gestures and their actions is what makes them memorable, as it is also suggested by Nacenta et al. [35]. Moreover, such memorable gestures do not only confirm users' perception of an easy-to-use device, but they also suggest that MemStone can be reliably used in infrequent settings, e.g., monthly meetings.

In general, most of our study participants expressed the desire to use MemStone in a context where devices are provided at the event location, e.g., meeting rooms (P2, P5, P13, P14, P15, P17, P19, P20). However, despite our prototype outperforming a comparable smartphone application, the interview sessions uncovered our participants' concerns for having to carry yet another device. Participants expressed their disinclination to carry a personal MemStone device with them during their everyday activities. While this suggests that convenience trumps efficiency [11], it might also be a as-of-yet too infrequent use case (controlling capture devices) to be of much use to people. A future filled with a plethora of capture situations may very well change this perspective.

As a possible limitation of our study, we acknowledge the fact that participants had to pretend that they are participating in a recorded meeting, which they watched through a laptop. However, even if we would have organized real meetings with our participants, they would also have to be scripted in one way or another, and hence would still be just as artificial as the recorded meetings. Nevertheless, we believe that evaluating the prototype in such lab scenarios still allowed participants to control different aspects of a meeting capture scenario. Results obtained from this study allow us to answer our research questions about the comparison of our interface with a more traditional smartphone interface, and explore participants' perceptions on our TUI prototype.

CONCLUSION AND FUTURE WORK

In this work we presented MemStone, a tangible user interface for controlling the capture and sharing of experience data in-situ. We evaluated our prototype in a meeting capture scenario with 20 participants. We found that our participants were significantly quicker in performing data capturing and sharing controls using MemStone than using a more conventional mobile app interface. The concept was highly valued by the participants, it was perceived as user-friendly, quick to learn, and easy and fun to use. Participants also expressed a positive attitude towards for the physical gestures and relationship with the control actions. We also found out that participants were able to remember the control gestures even after a long time period, which suggests that such TUI is suitable to be used in less frequently occurring events. However, in spite of its better performance and its high perceived value as an ambient-based control device, participants were very much divided about the convenience of having to carry an additional personal control device with them for their everyday activities.

We believe there is value in further improving our initial concept based on our participants' suggestions. One such improvement that we plan to address is to allow users to share their memories even with others that were not part of the same event. For example, by 'touching' two MemStones together, users could share data captured in the last hour, e.g., following a recent work by Geronimo et al. [12] on mid-air gestures.

ACKNOWLEDGMENTS

The authors acknowledge the financial support of the Future and Emerging Technologies (FET) programme within the 7th Framework Programme for Research of the European Commission, under FET Grant Number 612933 (RECALL). We would like to thank Matteo Pontiggia for his input in building a prototype of the MemStone interface, and Jacob Jeyaraj who helped in conducting the user study.

REFERENCES

1. Anne Adams. 2000. Multimedia Information Changes the Whole Privacy Ballgame. In *Proceedings of the Tenth Conference on Computers, Freedom and Privacy: Challenging the Assumptions (CFP '00)*. ACM, New York, NY, USA, 25–32. DOI: http://dx.doi.org/10.1145/332186.332199

2. Shane Ahern, Dean Eckles, Nathaniel S. Good, Simon King, Mor Naaman, and Rahul Nair. 2007. Over-Exposed?: Privacy Patterns and Considerations in Online and Mobile Photo Sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 357–366. DOI:http://dx.doi.org/10.1145/1240624.1240683

3. Claudio Bettini and Daniele Riboni. 2014. Privacy Protection in Pervasive Systems: State of the Art and Technical Challenges. *Pervasive and Mobile Computing* (Oct. 2014). DOI: http://dx.doi.org/10.1016/j.pmcj.2014.09.010

4. Agon Bexheti, Anton Fedosov, Jesper Findahl, Marc Langheinrich, and Evangelos Niforatos. 2015. Re-Live the Moment: Visualizing Run Experiences to Motivate Future Exercises. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15)*. ACM, New York, NY, USA, 986–993. DOI:http://dx.doi.org/10.1145/2786567.2794316

5. Agon Bexheti, Marc Langheinrich, and Sarah Clinch. 2016. Secure Personal Memory-Sharing with Co-Located People and Places. In *Proceedings of the 6th International Conference on the Internet of Things (IoT'16)*. ACM, New York, NY, USA, 73–81. DOI: http://dx.doi.org/10.1145/2991561.2991577

6. Simone Borsci, Stefano Federici, and Marco Lauriola. 2009. On the Dimensionality of the System Usability Scale: A Test of Alternative Measurement Models. *Cognitive Processing* 10, 3 (Aug. 2009), 193–197. DOI: http://dx.doi.org/10.1007/s10339-009-0268-9

7. Delphine Christin, Pablo Sánchez López, Andreas Reinhardt, Matthias Hollick, and Michaela Kauer. 2012. Privacy Bubbles: User-Centered Privacy Control for Mobile Content Sharing Applications. In *Information Security Theory and Practice. Security, Privacy and Trust in Computing Systems and Ambient Intelligent Ecosystems (Lecture Notes in Computer Science)*. Springer, Berlin, Heidelberg, 71–86. DOI: http://dx.doi.org/10.1007/978-3-642-30955-7_8

8. Sarah Clinch, Nigel Davies, Mateusz Mikusz, Paul Metzger, Marc Langheinrich, Albrecht Schmidt, and Geoff Ward. 2016. Collecting Shared Experiences through Lifelogging: Lessons Learned. *IEEE Pervasive Computing* 15, 1 (Jan. 2016), 58–67. DOI: http://dx.doi.org/10.1109/MPRV.2016.6

9. Sarah Clinch, Paul Metzger, and Nigel Davies. 2014. Lifelogging for 'Observer' View Memories: An Infrastructure Approach. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication (UbiComp'14 Adjunct)*. ACM, New York, NY, USA, 1397–1404. DOI: http://dx.doi.org/10.1145/2638728.2641721

10. Nigel Davies, Adrian Friday, Sarah Clinch, Corina Sas, Marc Langheinrich, Geoff Ward, and Albrecht Schmidt. 2015. Security and Privacy Implications of Pervasive Memory Augmentation. *IEEE Pervasive Computing* 14, 1 (Jan. 2015), 44–53. DOI: http://dx.doi.org/10.1109/MPRV.2015.13

11. Gurpreet Dhillon, Tiago Oliveira, Santa Susarapu, and Mário Caldeira. 2012. When Convenience Trumps Security: Defining Objectives for Security and Usability of Systems. In *Information Security and Privacy Research (IFIP Advances in Information and Communication Technology)*. Springer, Berlin, Heidelberg, 352–363. DOI: http://dx.doi.org/10.1007/978-3-642-30436-1_29

12. Linda Di Geronimo, Marica Bertarini, Julia Badertscher, Maria Husmann, and Moira C. Norrie. 2017. Exploiting Mid-Air Gestures to Share Data Among Devices. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*. ACM, New York, NY, USA, 35:1–35:11. DOI: http://dx.doi.org/10.1145/3098279.3098530

13. Tilman Dingler, Passant El Agroudy, Huy Viet Le, Albrecht Schmidt, Evangelos Niforatos, Agon Bexheti, and Marc Langheinrich. 2016. Multimedia Memory Cues for Augmenting Human Memory. *IEEE MultiMedia* 23, 2 (April 2016), 4–11. DOI: http://dx.doi.org/10.1109/MMUL.2016.31

14. Mingming Fan, Alexander Travis Adams, and Khai N. Truong. 2014. Public Restroom Detection on Mobile Phone via Active Probing. In *Proceedings of the 2014 ACM International Symposium on Wearable Computers (ISWC '14)*. ACM, New York, NY, USA, 27–34. DOI: http://dx.doi.org/10.1145/2634317.2634320

15. Pedro Felzenszwalb, Ross Girshick, David McAllester, and Deva Ramanan. 2010. Object Detection with Discriminatively Trained Part-Based Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 9 (Sept. 2010), 1627–1645. DOI: http://dx.doi.org/10.1109/TPAMI.2009.167

16. Md Sadek Ferdous, Soumyadeb Chowdhury, and Joemon M Jose. 2016. Privacy Threat Model in Lifelogging. ACM Press, 576–581. DOI: http://dx.doi.org/10.1145/2968219.2968324

17. Kenneth P. Fishkin, Anuj Gujar, Beverly L. Harrison, Thomas P. Moran, and Roy Want. 2000. Embodied User Interfaces for Really Direct Manipulation. *Commun. ACM* 43, 9 (Sept. 2000), 74–80. DOI: http://dx.doi.org/10.1145/348941.348998

18. George W. Fitzmaurice, Hiroshi Ishii, and William A. S. Buxton. 1995. Bricks: Laying the Foundations for Graspable User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '95)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 442–449. DOI: http://dx.doi.org/10.1145/223904.223964

19. Jim Gemmell, Lyndsay Williams, Ken Wood, Roger Lueder, and Gordon Bell. 2004. Passive Capture and Ensuing Issues for a Personal Lifetime Store. In *Proceedings of the the 1st ACM Workshop on Continuous Archival and Retrieval of Personal Experiences (CARPE'04)*. ACM, New York, NY, USA, 48–55. DOI: http://dx.doi.org/10.1145/1026653.1026660

20. Daniela Grijincu, Miguel A. Nacenta, and Per Ola Kristensson. 2014. User-Defined Interface Gestures: Dataset and Analysis. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces (ITS '14)*. ACM, New York, NY, USA, 25–34. DOI:http://dx.doi.org/10.1145/2669485.2669511

21. Morgan Harvey, Marc Langheinrich, and Geoff Ward. 2016. Remembering through Lifelogging: A Survey of Human Memory Augmentation. *Pervasive and Mobile Computing* 27 (2016), 14–26. DOI: http://dx.doi.org/10.1016/j.pmcj.2015.12.002

22. Lars Erik Holmquist, Friedemann Mattern, Bernt Schiele, Petteri Alahuhta, Michael Beigl, and Hans-W. Gellersen. 2001. Smart-Its Friends: A Technique for Users to Easily Establish Connections between Smart Artefacts. In *Ubicomp 2001: Ubiquitous Computing*. Springer, 116–122. DOI: http://dx.doi.org/10.1007/3-540-45427-6_10

23. Roberto Hoyle, Robert Templeman, Denise Anthony, David Crandall, and Apu Kapadia. 2015. Sensitive Lifelogs: A Privacy Analysis of Photos from Wearable Cameras. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1645–1648. DOI: http://dx.doi.org/10.1145/2702123.2702183

24. Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014. Privacy Behaviors of Lifeloggers Using Wearable Cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. ACM, New York, NY, USA, 571–582. DOI:http://dx.doi.org/10.1145/2632048.2632079

25. Hiroshi Ishii and Brygg Ullmer. 1997. Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI '97)*. ACM, New York, NY, USA, 234–241. DOI: http://dx.doi.org/10.1145/258549.258715

26. Apu Kapadia, Tristan Henderson, Jeffrey J. Fielding, and David Kotz. 2007. Virtual Walls: Protecting Digital Privacy in Pervasive Environments. In *Pervasive Computing (Lecture Notes in Computer Science)*. Springer, Berlin, Heidelberg, 162–179. DOI: http://dx.doi.org/10.1007/978-3-540-72037-9_10

27. Marion Koelle, Wilko Heuten, and Susanne Boll. 2017. Are You Hiding It?: Usage Habits of Lifelogging Camera Wearers. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*. ACM, New York, NY, USA, 80:1–80:8. DOI: http://dx.doi.org/10.1145/3098279.3122123

28. Mohammed Korayem, Robert Templeman, Dennis Chen, David Crandall, and Apu Kapadia. 2016. Enhancing Lifelogging Privacy by Detecting Screens. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 4309–4314. DOI: http://dx.doi.org/10.1145/2858036.2858417

29. Kai Kunze, Masakazu Iwamura, Koichi Kise, Seiichi Uchida, and Shinichiro Omachi. 2013. Activity Recognition for the Mind: Toward a Cognitive "Quantified Self". *Computer* 46, 10 (Oct. 2013), 105–108. DOI:http://dx.doi.org/10.1109/MC.2013.339

30. James R. Lewis and Jeff Sauro. 2009. The Factor Structure of the System Usability Scale. In *Human Centered Design (Lecture Notes in Computer Science)*. Springer, Berlin, Heidelberg, 94–103. DOI: http://dx.doi.org/10.1007/978-3-642-02806-9_12

31. Li-Jia Li and Li Fei-Fei. 2007. What, Where and Who? Classifying Events by Scene and Object Recognition. In *2007 IEEE 11th International Conference on Computer Vision*. 1–8. DOI: http://dx.doi.org/10.1109/ICCV.2007.4408872

32. Allan Christian Long, Jr., James A. Landay, and Lawrence A. Rowe. 1999. Implications for a Gesture Design Tool. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*. ACM, New York, NY, USA, 40–47. DOI: http://dx.doi.org/10.1145/302979.302985

33. Matthew B. Miles and A. Michael Huberman. 1994. *Qualitative Data Analysis: An Expanded Sourcebook*. SAGE.

34. Simon Moncrieff, Svetha Venkatesh, and Geoff West. 2008. Dynamic Privacy Assessment in a Smart House Environment Using Multimodal Sensing. *ACM Trans. Multimedia Comput. Commun. Appl.* 5, 2 (Nov. 2008), 10:1–10:29. DOI: http://dx.doi.org/10.1145/1413862.1413863

35. Miguel A. Nacenta, Yemliha Kamber, Yizhou Qiang, and Per Ola Kristensson. 2013. Memorability of Pre-Designed and User-Defined Gesture Sets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 1099–1108. DOI: http://dx.doi.org/10.1145/2470654.2466142

36. Donald A. Norman. 1999. Affordance, Conventions, and Design. *interactions* 6, 3 (May 1999), 38–43. DOI: http://dx.doi.org/10.1145/301153.301168

37. Donald A. Norman. 2007. *Emotional Design: Why We Love (or Hate) Everyday Things*. Hachette UK.

38. Kieron O'Hara, Mischa M. Tuffield, and Nigel Shadbolt. 2009. Lifelogging: Privacy and Empowerment with Memories for Life. *Identity in the Information Society* 1, 1 (Feb. 2009), 155–172. DOI: http://dx.doi.org/10.1007/s12394-009-0008-4

39. Blaine A. Price, Avelie Stuart, Gul Calikli, Ciaran Mccormick, Vikram Mehta, Luke Hutton, Arosha K. Bandara, Mark Levine, and Bashar Nuseibeh. 2017. Logging You, Logging Me: A Replicable Study of Privacy and Sharing Behaviour in Groups of Visual Lifeloggers. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 2 (June 2017), 1–18. DOI:http://dx.doi.org/10.1145/3090087

40. Jun Rekimoto and Eduardo Sciammarella. 2000. ToolStone: Effective Use of the Physical Manipulation Vocabularies of Input Devices. In *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology (UIST '00)*. ACM, New York, NY, USA, 109–117. DOI:http://dx.doi.org/10.1145/354401.354421

41. Jennifer Sheridan, Ben W. Short, Kristof Van Laerhoven, Nicolas Villar, and Gerd Kortuem. 2003. Exploring Cube Affordance: Towards a Classification of Non-Verbal Dynamics of Physical Interfaces for Wearable Computing, Vol. 2003. IEE, 113–118. DOI: http://dx.doi.org/10.1049/ic:20030156

42. Michael E. Sobel. 1982. Asymptotic Confidence Intervals for Indirect Effects in Structural Equation Models. *Sociological Methodology* 13 (1982), 290–312. DOI: http://dx.doi.org/10.2307/270723

43. Robert Templeman, Mohammed Korayem, David Crandall, and Apu Kapadia. 2014. PlaceAvoider: Steering First-Person Cameras Away from Sensitive Spaces. In *Network and Distributed System Security Symposium (NDSS)*. DOI: http://dx.doi.org/10.14722/ndss.2014.23014

44. Elise van den Hoven, Joep Frens, Dima Aliakseyeu, Jean-Bernard Martens, Kees Overbeeke, and Peter Peters. 2007. Design Research & Tangible Interaction. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction (TEI '07)*. ACM, New York, NY, USA, 109–115. DOI: http://dx.doi.org/10.1145/1226969.1226993

45. Kristof Van Laerhoven, Nicolas Villar, Albrecht Schmidt, Gerd Kortuem, and Hans Gellersen. 2003. Using an Autonomous Cube for Basic Navigation and Input. In *Proceedings of the 5th International Conference on Multimodal Interfaces (ICMI '03)*. ACM, New York, NY, USA, 203–210. DOI: http://dx.doi.org/10.1145/958432.958472

46. Tai-Kuei Yu, Long-Chuan Lu, and Tsai-Feng Liu. 2010. Exploring Factors That Influence Knowledge Sharing Behavior via Weblogs. *Computers in Human Behavior* 26, 1 (Jan. 2010), 32–41. DOI: http://dx.doi.org/10.1016/j.chb.2009.08.002