Countertop Responsive Mirror: Supporting Physical Retail Shopping for Sellers, Buyers and Companions

Maurice Chu, Brinda Dalal, Alan Walendowski and **Bo Begole**
Palo Alto Research Center, Inc. (PARC)
3333 Coyote Hill Road, Palo Alto, CA 94304, <first.last>@parc.com

ABSTRACT
We examine opportunities for ubiquitous technologies in retail shopping, jewelry shopping in this case, to supplement the unique information needs inherent to physical trials of tactile products. We describe an iterative design approach to develop a mirror system that records and matches images across jewelry trials called the Countertop Responsive Mirror. The key technological distinction of our system from prior technologies is the use of "matched access," which automatically retrieves images that match a scene shown in separately accessed images. This not only helps shoppers compare jewelry but also promotes interactions among all parties during shopping. We report qualitative findings from multiple field trials of the system. This paper contributes to a body of research on the design and introduction of new technologies into retail shopping that provide value to all users without disruption to their normative practices and behaviors.

Author Keywords
Natural and implicit interaction, retail, camera system, computer vision application, digital mirror, shopping.

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
Making a purchase decision involves gathering and assessing multiple elements of information like alternatives, features, reviews, availability, cost, size, colors, etc. Such information needs underlie the success and continued expansion of information technologies for online shopping. Using the web, shoppers can find a great deal of information about products they are interested in and they can search and compare alternative choices side by side.

Although online shopping is used across many product categories, the assessment of products in certain categories, such as furniture, apparel, eyeglass frames and jewelry products, usually includes some form of tactile and physical information: texture, fit, drape, flow, movement, light reflection, heft, etc. These kinds of information are difficult to communicate electronically because they use human sensing modalities that are not easily quantifiable for electronic transfer and/or are based on each individual's subjective perception. Underhill has elaborated on these and other factors that continue to compel shopping in physical retail stores [25].

In addition to these tactile properties, there are certain intangible aspects that determine how well the article *fits* the shopper. Fit is determined not only by the dimensions of the product, but also by how well suited the item is to the *presentation of self* that the shopper intends to project, sometimes referred to as "style". Product selection also involves assessing information about the social context of products like style trends and the personal opinions of others, typically family and friends [6, 15, 17 and 18].

So far, the only technologies we see in widespread use on the floor of a retail store are ones that reduce the need for sales staff (price-check scanners, buttons to request assistance, automated assistance kiosks and marketing displays). A number of research prototypes aimed to aid people's in-store experience are based on mobile phone applications that bring web-based information into the store (e.g. [20]), sometimes using RFID tags, optical codes [19] or visual matching engine to match photos of objects in a store against a database of similar objects [26]. We do not see many technologies designed to augment the unique information needs inherent to physical trials of tactile products while not disrupting the normative processes of sellers, buyers and their companions. Such design goals are the aim of the research reported in this paper.

BUYER AND SELLER CO-DEPENDENCE IN RETAIL
Often, the relationship between salespeople and shoppers is perceived as predatory, even parasitic, where sales people manipulate buyers into decisions that are perhaps not wholly in the buyer's best interest. This antagonistic perspective of the buyer-seller relationship stems from competing goals in the sales transaction: the buyer wants to spend as little as possible and the seller wants to acquire as much as possible for the merchandise.

While these competitive goals underlie their interactions, the buyer and seller also share an objective: satisfying the buyer's criteria with available merchandise. The buyer and seller have complementary knowledge to bring to bear on

CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.

Figure 1. Cover of Spring 2004 Shaadi Style Magazine. (image used under "fair use")

reaching a resolution. The buyer has at least a vague notion of what she wants but not the full range of a seller's options; the seller knows what is available but not what would optimally fit the shopper's criteria. The sales process is a negotiation between these parties to optimize their mutual objective. From this perspective, shopping is a collaborative information seeking process with the ultimate goal of making a mutually satisfying sales transaction.

Sales practices have been examined in marketing and sociological research. In *Making Sales*, Robert Prus [21] describes the main elements of the sales process.

1. Promote interest
2. Generate trust
3. Neutralize reservations
4. Obtain commitments (to purchase, lease, try, etc.)
5. Manage disruptions from others
6. Foster long-term relationships

For buyers, many factors influence their purchase decisions, some of which are difficult to quantify, such as how shoppers use their selections to foster social relationships as described by Miller [17, 18]. Though not well understood, consumer behaviors receive considerable coverage in sales and marketing texts [4, 7, 11 and 25]. Most describe a process in which consumers evaluate product alternatives by weighing multiple criteria and then applying a variety of decision rules to make a choice. Evaluation criteria range from basic attributes of the product, to attitudes about brand and lifestyle, as well as affective factors and emotional response. To receive effective assistance from a seller, the buyer must communicate their criteria and weightings to the seller while simultaneously withholding information that could result in paying more than desired.

South Asian Retail Jewelry

We chose to study South Asian retail jewelry because it is a $13 billion dollar industry (with projected annual growth of 8.7% through 2025 [5, 16]). If our designs are successful, this will be a substantial market in which to have impact.

Apart from being decorative, jewelry is often purchased in India for special occasions: annual spring and autumn festivals, to celebrate the birth of newborns, marriages and birthdays and the purchase of 22 and 24 carat gold and diamonds are also seen as a long term investment. Indian jewelry ensembles are often elaborate combinations of large necklaces with matching earpieces along with a variety of bangles, rings, forehead pieces and more. Figure 1 shows a magazine cover with a photograph of a woman in a typically elaborate ensemble. Shoppers spend considerable time examining and selecting such items before purchasing.

To understand the current physical layout and practices in Indian retail jewelry, we informally interviewed more than a dozen people who have shopped in jewelry stores in India (including one member of the research team) and visited four stores in the San Francisco bay area that specialize in Indian jewelry. We interviewed 3 of those store owners and at least one salesperson from each. We observed the process in the stores for more than 20 hours. We also examined photographs of jewelry stores in India looking specifically at store layouts and jewelry placement. Although we observed a small but prominent representative subset of South Asian jewelry stores, the following descriptions do not necessarily characterize all Indian jewelry shopping.

Jewelry Store Layout

Based on our sources, a typical Indian jewelry store is laid out similarly to jewelry stores in other parts of the world but with a few significant differences. Jewelry items are on display often laid out on velvet trays, either in glass counters which are about mid-torso height and run the circumference of the room, or decoratively placed on shelves on the walls. Items are often grouped by type; with necklaces, bands, rings and other pieces clustered together. Several small, portable mirrors sit on glass counters, and large mirrors are installed on the walls as well. A significant variation from other types of jewelry store layout is that matching sets of elaborate neck and earpieces are often placed prominently at eye level, on the wall.

The setup of the counters allows shoppers to easily browse all the jewelry items while keeping the jewelry items secure from theft. Also, salespeople are able to interact with shoppers in close proximity since they can follow the shoppers as they walk around the store. The grouping of jewelry items makes it easier for shoppers to compare similar items side-by-side.

Observed Shopping Process

The buyer-seller relationship in Indian jewelry stores is even more important than in western-style retail. In India, families buy from the same jeweler across multiple generations. Furthermore, shoppers expect in all cases to be served by a salesperson, not left to simply browse on their own. Generally, an Indian jewelry store will have enough sales staff to serve each shopper individually. This strong buyer-seller relationship provides an acute situation where disruptions would quickly become apparent.

When shoppers enter a jewelry store, they browse the store's inventory by looking through the counters. Salespeople assist customers by standing behind the counters across from the customers. The salesperson will retrieve an item of interest and hand it to the shopper to try on. If a mirror is not nearby, the salesperson will bring the nearest mirror over. The shopper then tries on the jewelry and looks in the countertop mirror for close-up views and the mirrors on the walls for evaluating the whole look.

Shoppers usually come with their friends or family, in larger groups than typical in western jewelry stores, so the process is a social one involving looking at the mirror, posing for companions and exchanging feedback. Shoppers sometimes get feedback from the salespeople who use this time as an opportunity to suggest similar items to try on. Also, shoppers will bring in clothing for which they wish to buy matching jewelry so that they can see the jewelry in the context of the clothing they will be wearing. When the shopper asks to try on other items, the salesperson takes care to monitor the jewelry items and will often place jewelry items back into the glass counters or on the shelves.

After trying on several items, the shopper may ask to see a few of them again. The salesperson will retrieve those items and place them on a tray in front of the shopper. Many Indian jewelry items are larger than western counterparts, and the tray is compartmented into 5 or 6 segments to hold the matching elements of ensembles. The shopper then compares the jewelry items against each other with far more scrutiny than before and tries the items on again if desired. A mirror is brought over to help the shopper make a decision, and the shopper's friends and families give feedback to help evaluate the jewelry in more detail.

Jewelry Shopping Tools

During our onsite observations, we noted the prevalence of three recurring artifacts: mirror, tray and weight scales. The most prominent tool that directly contributes to shoppers' product evaluation is the mirror. Interestingly, the mirror is also a common point of interaction between the buyer and seller, and also between the buyer and her companions. Sellers typically arrange the customer's preferred jewelry pieces on a light and portable tray so that the shopper can compare items side by side. Many items have fixed prices attached to them but for some items the salesperson weighs it on a small scale to determine its price according to the daily gold and silver prices.

The mirror and tray appeared to be the most natural place to insert technology that can both enable better collaborative information seeking and minimally disrupt current shopping practices. The weight scale, on the other hand, is used only by the seller. There is a marked lack of interaction between the buyer and seller whenever the scale is employed. Because our goal is to supplement, not disrupt, the normative collaborative practices, the scale did not provide an opportunity to enhance the interaction. With regard to the mirror and tray, we observed the following practices.

- Shoppers use multiple criteria when evaluating jewelry products and they have been observed to try on jewelry several times as they assess different aspects.
- Mirrors provide an opportunity for salespeople to give feedback to the shopper and recommendations.
- The mirrors and trays are also helpful in engaging the shopper's companions, as friends and family will suggest other items to try on and even hold up the jewelry to the shopper's neck and ears to get an idea of how the jewelry would look without the shopper having to put it on.
- During the final decision to buy, trays containing the top selected jewelry items are provided to the shopper for detailed evaluation and comparison. In order to see the jewelry items as it fits on the shopper, she must try the jewelry on again but usually does not.

COUNTERTOP RESPONSIVE MIRROR SYSTEM

We followed an iterative design process in the development of our Countertop Responsive Mirror system. The version we describe in this section is the end result of four cycles of deployments in real jewelry stores and is referred to as P2.

System Hardware

We have developed the system with the dual goals of (1) enabling shoppers to better evaluate jewelry and (2) serve as a platform where storeowners, shoppers, and their companions can interact with one other. The prototype consists of two components: (Figure 2) a mirror for trying on and recording items and (Figure 3) a display for

Figure 2. Countertop Responsive Mirror (P2). Mirror component consists of camera behind half-silvered mirror.

reviewing and comparing items. This separation of function was a deliberate design decision to match the normative shopping practices of using the conventional mirror and tray for the two phases of shopping.

The recording mirror is meant to take the place of the conventional mirrors found in jewelry stores, with the added ability to record images and video. The mirror component (Figure 2) consists of a half-silvered mirror with an embedded camera and an embedded 8 inch color Liquid Crystal Display monitor which showed live feedback of the camera view, providing peripheral awareness of the recording without requiring them to look away from the mirror. The associated silver-colored PowerMate knob serves as a "record/stop recording" button that controls the camera capture. We intentionally designed it for portability and individual use and the mirror was sized accordingly (8 x 10 inches, 20 x 25 centimeters), which is similar to the sizes of small portable mirrors we observed in stores. This size causes the shopper's head and neck to fill most of the image, reducing background images that can impede accuracy of the computer vision algorithms.

The comparison component is a touch screen display with a GUI that enables shoppers to view and compare their recorded sessions (Figure 3). It is intended to take the place of the tray used for reviewing top selections and allow the shopper to see the items as worn. We also developed an image matching algorithm, which matches images that have the same body/head pose, to implement our "matched access" to compare two recorded videos. This is a distinguishing feature of our system that enables shoppers to quickly compare two or more recordings of themselves from any orientation, like front view and side profile, wearing different jewelry in a variety of poses. The reviewing component is meant for social use to support interactions among shoppers, salespeople, and the shoppers' companions. Users can interact with the touch screen to play back images, using a slider bar to view different frames. A touch screen proved to be more suitable than a mouse and allowed multiple people to manipulate the viewer together.

The PC is hidden from sight, and shoppers see only the mirror with a large PowerMate push button in front of it. The touch screen display may be placed off to the side or in a wholly separate location, allowing another shopper to use the mirror component while the first shopper reviews items.

System Software

The software has three logical components: the GUI, in managed C++; the session management layer, a native C++ interface on top of MySQL; and the image manipulation code which is also in native C++ and developed on top of OpenCV [3]. When idle, both the embedded display and the touch screen play an introductory video demonstrating use of the Countertop Responsive Mirror and inviting the shopper to press the button to begin.

To begin, a shopper must identify herself by logging in or creating an account. For our trials, we simply created anonymous accounts at each login. In a later prototype (not yet deployed), each shopper is given a card with a unique optical code that they hold in front of the camera to identify their personal sets of images. They can choose to associate a personal account with that identifier or not, either in the store or later from a web site.

Once logged in, the shopper simply uses the system as she would a normal jewelry mirror, except to press the PowerMate button before and after trying on a new piece of jewelry to start and stop the acquisition of a sequence of images[1]. When recording, the Acquisition Engine displays the text "RECORDING" and a live preview is shown of the camera's field of view, showing the images being captured on the embedded LCD monitor in the mirror. Rather than saving every single frame that is captured by the camera, the Image Matching Engine is invoked to discard frames that are not different enough from previous frames. After recording has finished, the Acquisition Engine saves the newly acquired session of images.

Once the shopper has recorded an item of jewelry, a thumbnail of the first image recorded in the session appears on the display in the mirror as well as on the touch screen to provide visual feedback that the recording has been saved. Each new recording session creates a new thumbnail, allowing the user to see the previously captured items at a glance (see bottom of Figure 3). The thumbnails may be dragged to the trash icon located on the bottom left of the touch screen to delete the session or may be dragged into either of the two larger panels above for comparison.

The comparison panels display individual frames of whichever sessions are dragged in to them. Once sessions are dragged into the panels, the shopper can play back the saved frames of a session by pressing the Play button, or by

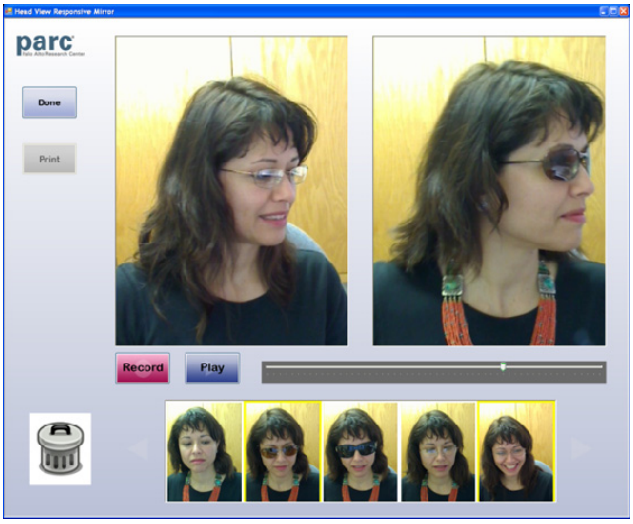

Figure 3. Recordings are compared using the above interface on a touch-screen "review" component

dragging a slider below the comparison panels. While playing, the images in the panels appear video-like. Shoppers can use the slider to scan the images in a session and select particular images to view in detail.

When both comparison panels are loaded with two different sessions, the Image Matching Engine algorithmically matches each image in one session to the image in the other session which has the most similar pose of the shopper. For example, in Figure 4, the column of images on the left displays one session where the customer is trying on a long necklace. Accordingly, in the right hand column, the Image Matching engine has selected and displayed poses from a second session which most closely match those of the customer looking to the left, right, and straight on. Note that the images in the second panel do not necessarily show the entire set of captured images from the second session and are not presented in the same sequence in which they were acquired. This *"matched access"* capability enables shoppers to compare how two items (or more) look in similar poses.

Matching Algorithm

Rather than using computer vision algorithms to estimate the pose of the person by tracking the head, hands, arms, and torso [9], we chose to take a fresh approach by engineering a similarity measure between two images, which corresponds to what people would perceive as the best pose-matched image across sessions.

We started by considering the sum-of-squared (SSD) distances between pixels of two RGB color images, I_1 andI_2, as a baseline. We found that this simple distance measure worked surprisingly well when there were only 10-20 images in each session, and it was even able to handle cluttered and moving backgrounds, like family and friends in the field of view, because the shopper takes up most of the field of view (and hence, most of the pixels in the image). One case that was problematic was that slight translations of the shopper caused large changes in the SSD distance, whereas human perception is less sensitive to translations in body position. To make the distance measure less sensitive to translations and because we expect that most pixels will be coming from the shopper and not other background clutter, we modified the SSD distance to allow each pixel in the first image to either be associated with the same pixel in the other image or be shifted by a common translation (dx, dy):

$$d_{SSDmod}(I_1, I_2) = \min_{dx,dy} \sum_{x,y \in P} \min \begin{Bmatrix} \|I_1(x,y) - I_2(x,y)\|^2, \\ \|I_1(x,y) - I_2(x+dx, y+dy)\|^2 \end{Bmatrix}$$

where P is the set of pixel coordinates. The final implementation includes optimizations to allow for numbers of images per session to be on the order of 100-150 with adequate responsiveness.

Figure 4. Matching Example.

Regarding matching performance, we require that pairs of images are "close enough" for a user to perform a visual side-by-side comparison. We collected a dataset using our prototype of four participants each recording 3-5 sessions, each consisting of 48 to 151 images, wearing different necklaces. We found that the most salient pose variations in the captured images were the head orientation (left (L), center (C), right (R)). We hand labeled each image in our dataset to use as an estimate of "ground truth." After running our matching algorithm for each pair of sessions within participants, we generated a confusion matrix of the image label of the reference image with the image label of the computed matched image.

Table 1. Confusion Matrix of Head Orientation

Reference image pose	Matched image pose		
	L	C	R
L	110 (2.8%)	402 (10.2%)	89 (2.3%)
C	381 (9.7%)	1616 (41.0%)	437 (11.1%)
R	89 (2.3%)	446 (11.3%)	368 (9.3%)

The head orientation is matched precisely only 53% of the time (which is the diagonal of black entries in the confusion matrix). However, matching a left-oriented head image to a center-oriented head image is not perceived as being as poor a match as matching a left-oriented head image to a right-oriented one. The head orientation is nearly matched correctly 95% of the time (which is the sum of the black and gray entries in the confusion matrix). In our deployments described later, users did not make note of cases when the system had a near mismatch but sometimes

did call our attention to a large mismatch (4.6%, the sum of white entries in the confusion matrix). Users no doubt perceived many of the near misses, but from their reactions such misses were not as severe as the large misses.

The above approach is hardly a technical contribution to computer vision; however, the contribution of the above method is to demonstrate that simple vision methods can be sufficient, depending on the constraints of an application. The above method does not extract faces and body parts, much less extract pose parameters, but it can match poses in images including head turn, leaning forward and use of hands in the image which fits the needs of this application.

RELATED TECHNOLOGIES

Capture and Access Systems

Capture and access systems [23] are a class of ubiquitous computing systems that capture parts of an experience, via interactions with a user interface, cameras, and microphones, for access later. Several systems from note taking in classrooms, recording meal preparation in the home, capturing informal meetings at work and battlefield visualizations in the military domain have been developed. One system called the cook's collage [24] has been developed as a short-term memory aid when interrupted while in the middle of cooking to help recover what the cook had been doing before being interrupted.

The Countertop Responsive Mirror falls into the class of capture and access systems, serving as a memory aid of prior events. A contribution of this system is an extension to the basic notion of capture and access to that of "matched access" for comparing between captured sessions. "Capture and access" involves recording an event and creating an index that is used to access the recording. In prior art, the "capture" might use technology to auto-record or auto-generate the index but users explicitly access the content. The "matched access" of this system is the first to our knowledge that uses technology to automatically access a captured event, matching it to a scene in the real world or a scene in a separately controlled playback.

Collaborative Face to Face Consultation Systems

The work by Rodden et al. [22] has focused on the design of information technology for high-value, complex products like hi-fi systems, financial portfolios and travel packages where a face-to-face consultation with an expert agent is often used to customize the product to what the customer wants and is able to purchase. They redesigned the physical layout of the displays, changed the seating arrangement to be a shoulder-to-shoulder one, and implemented a visual interface for both the agent and customer to result in a much more collaborative experience.

Off-the-shelf/rack/counter retail apparel shopping, which our Countertop Responsive Mirror system addresses, is for products that may be high-value and complex, but are not highly customizable. When buying clothing, jewelry, glasses, and other apparel, customer choice is largely confined to selecting among the decisions made by fashion designers before the products reach the store. Unlike, travel and financial products that generally involve an agent experienced with the range of available choices to assist in configuring a package, in retail apparel sales, shoppers can generally see for themselves what the available choices are and the interaction between salespeople and shoppers involves some assistance in accessing back-room inventory (for alternate size, color, etc.) and also in helping the customer assess the *suitability* of the product which includes aspects of the product's fit and style.

Interactive Mirrors

In 2001, Prada opened a flagship store in Manhattan, New York where it conducted a trial of many technologies, including a sophisticated dressing room with a variety of capabilities [14]. A scanner identified each garment as the shopper took it into the room, providing additional information about the garment's price and alternate colors and sizes – the same kind of information shoppers can find when browsing products online. The fitting room also contained a Magic Mirror with a motion-triggered video camera that recorded the shopper and played back the video after a pause. The system also provided the ability for a person trying on clothes to send video of himself to friends who can send back comments and vote (thumbs up/down). The system could also project a static image of an alternate garment onto the mirror providing a basic "virtual fitting" capability with which the shopper could get some sense of how the garment might look on him. The trials of these technologies were not successful. A report in *Business 2.0* describes the dramatic mismatch between expectations of the retail technology designers and the reality of use of the technologies day to day where much of the system went unused due to a variety of factors including overflow traffic, technical failures and non-intuitive controllers (such as floor pedals to set the opacity of a glass wall) [14].

Prior research has investigated the design and use of a Responsive Mirror for trying on *clothing* [2 and 27], rather than jewelry which is the focus of the currently described work. The previous system used two cameras, a front camera to capture images and recognize classes of clothing, and an overhead camera to detect orientation toward the mirror, whereas the system we describe in this paper requires only a single camera embedded behind a half-silvered mirror that captures frontal images of shoppers and uses a different vision algorithm to match images with similar head and body pose. Although formative trials of the system indicated some promise, the introduction of cameras into the semi-private space of evaluating clothing (though not the actual changing of clothing) is a significant hurdle for adoption of the technology [2]. The system described here is focused on a more public shopping experience which mitigates the privacy implications.

The mirror metaphor has been employed in the domains of cosmetics and also the trial of eyeglass frames. The Smart

Makeup Mirror [12] uses a high-resolution camera and monitor to provide functionality analogous to a digitally enhanced lighted dressing-table magnifying mirror. The user can zoom into specific regions of the face and see how colors change in simulated lighting conditions. Commercially, ACEP sells a product called Smart Mirror [1] consisting of a camera that takes still images of a shopper wearing different eyeglass frames and displays them side-by-side. Smart Mirror does not provide the pose matching of the Countertop Responsive Mirror.

The Philips MiraVision LCD Mirror TV is an example of a commercial product that integrates an electronic display *behind* a half-silvered mirror. There are some research systems exploring techniques to optimize the position of what is shown in the electronic display with what is reflected by the mirror [10 and 13]. These systems do not perform electronic matching of images.

FIELD TRIALS

This section highlights the results of our field trials. The trial deployments aimed to uncover the affordances of artifacts used in jewelry stores and to identify shopping practices and customer and sales interactions critical to the design of prototypes appropriate for the South Asian market. We conducted these trials not to test any specific hypotheses, but to identify the high-order issues affecting the usefulness of the system to all user types (sellers, buyers and companions. Lengthy installations for quantitative tests were not possible because it is too costly to deploy and support prototype iterations over extended periods in an operating store. Thus, what we learned first-hand from several shorter trials were the realistic constraints and opportunities of deploying technology in jewelry stores where space, time and resources are limited.

In-Store Trials

We deployed versions of the prototype for three periods in two jewelry stores. An engineer and a social scientist made independent observations and took detailed field notes while on the floor. Specifically, we tracked the number of people that came into the store, whether they were alone, in pairs, or in groups and how long they spent in the store; their interactions with artifacts, including timing how long customers spent in front of mirrors, and also made thorough field notes of conversations and actions between sales people and customers.

Table 2. Categories of shoppers observed in store deployments.

Number of	Shop A (23 shoppers)	Shop B (24 shoppers)	Total (47 shoppers)
Units	12	13	25
People who tried on jewelry	7 (58%)	8 (62%)	15 (60%)
People who purchased items	2 (16%)	5 (38%)	7 (28%)
People who used CRM	5 (42%)	5 (38%)	10 (40%)

We observed a total number of 47 shoppers, specifically, 23 shoppers in Shop A and 24 shoppers in Shop B, summarized in Table 2. Overall, 29 were female and 18 male (including children). Across these, we identified three primary *units* consisting of individual shoppers, pairs of shoppers and people in groups of three or more. The total number of units amounted to 25. Only five units contained children. Of the 12 units in Shop A, the majority of people observed were in pairs (7 units), or alone (4 units), with only one family group. The majority of shoppers in Shop B consisted of individuals (6 units), with four pairs and three groups. Shoppers spent between 40 minutes to two minutes in the store, averaging out, overall, to about 10 minutes per unit. The following table displays the percentage of units of people who tried on and purchased jewelry and used the Countertop Responsive Mirror.

Altogether, 15 people used the prototype. Of these, 10 were observed in the stores and 5 filled in a survey without being observed by the study team. In Shop A, the store manager explicitly encouraged the use of the prototype whereas the owners of Shop B did not, preferring that sales people not attempt to learn or use the technology.

The first version of the prototype (P1), installed for three days at Shop A, a higher end jewelry store, consisted of a free standing camera that could be attached to any pre-existing mirror and a PC, PowerMate knob and monitor to capture and play back recorded images. We attached the camera at eye level to a large mirror in the middle of the store some distance away from the jewelry counters, and placed the monitor on a table nearby. The store manager controlled the recordings (captured through the web cam) and playback sessions (displayed on the monitor). .

In addition to direct observations in-situ, the research team also created a survey with 15 questions asking customers to describe their experience with the CRM, such as whether they found it useful to compare different images, whether it saved them time, what features they found easy or hard to use, and whether they preferred still images or video recordings of themselves. Five shoppers (who were not observed) filled out the survey. Four out of five rated the image comparison feature highly useful (5 on a 5 point scale) while one thought it was fairly useful (3). Four people stated that the tool would help them to decide on what to buy. As one remarked, it provided her with the "ability to go back and forth [between the different images] without having to try pieces on over and over again". Another commented that the tool "helps me choose [between items]" and that "comparing [across items] was easy". Respondents also suggested improving the intuitiveness of the user interface regarding consistency across buttons and instructions about recordings, as one explained, "When I stand in front of the equipment, it's not obvious what I need to do."

Orienting to the Camera

The manager of Store A was briefly trained in how to use the technology, including recordings of customers trying on the jewelry (with their permission), playing back recordings and generally managing the computer. He was highly enthusiastic about the technology and found it easy to use in terms of handling the buttons on the touch screen, recording images and dragging thumbnails to a preview pane. Having salespeople control the device was initially perceived to be more realistic for India. For customers, however, we observed that the arrangement was perplexing because they did not know how they would appear in the image. Shoppers couldn't adjust their jewelry, hair, clothing or proximity to the camera and were unable to linger on the images as they might have done while trying on jewelry in front of a conventional mirror. We noticed that salespeople rarely had the time to ensure that the images of the jewelry and shoppers were displayed at the best angle or in the best light. Backlit, shadowy images and unflattering recordings caused people to quickly lose interest. Shoppers found it disconcerting to alternate their gaze between the recording device at eye level and the playback device at waist level.

Unlike retail clothing stores where people can gather up and try on an armful of clothes, in jewelry stores, shoppers tended to select one or two items to try on at a time, and browsed the counters between selections. We found that the placement of conventional mirrors, distributed across several counters, helped to maintain a sense of privacy and intimacy while people sampled jewelry in a public space. Thus when our web camera was originally attached to a large, in-store mirror, shoppers were self-conscious of having to walk across the store to try on their jewelry in full view of other shoppers.

A related phenomenon observed among *all* our participants who tried on necklaces or earrings was a repetitive body movement towards and away from the mirror. After putting on a piece of jewelry, people leaned in towards the mirror, their faces a few inches away from its surface, to see close-up the jewelry that they had tried on. They would then straighten up or take a step back so that they could see their head, neck and shoulders in the mirror, and adjust their clothing or hair to better view the necklace against their collarbone, or against the neckline of their clothing.

Based on these observations, we realized that it was not sufficient to simply place a web camera on an existing, large mirror in the store. The variety of observed behaviors made it clear that restricting the poses that the system could handle would be detrimental to the shopper's needs, so that matching poses in recorded images must be able to handle any pose that the shopper wishes to view herself in.

Thus we chose to redesign the prototype. In the second prototype (P2), we deliberately constrained the width of the countertop responsive mirror. By embedding a camera within the mirror, we could also more effectively capture people's natural actions.

During the second deployment weeks later in the same jewelry store, shoppers responded more favorably to the revised prototype (P2) after we had embedded the camera and visual feedback capability in the mirror. The placement of the prototype was changed to be nearer to a jewelry counter where it blended in, rather than stood apart and the sales manager could invite customers to try it.

Customers wished to have images available beyond store boundaries in order to obtain family members' opinions about items, especially if their spouse was unable to accompany shoppers to the store but were expected to purchase items. In contrast, store managers, cautious about designs being shared and copied by competitors, did not respond favorably to the notion that customers could send images of proprietary designs outside the store.

Our third deployment was of P2 at a family-owned jewelry store, where the store owners were less familiar with computer technology and were not comfortable in handling the device, except to reboot the PC every morning. Here, 62% of people observed in the store tried on jewelry and 38% of them purchased items. 38% also interacted with the Countertop Responsive Mirror, which largely consisted of spontaneously watching the looped video. In one case, a woman who entered the store saying that she wanted to buy a set of gold earrings for someone as a gift ended up buying two sets of diamond studs after using the CRM to try on several pairs.

In contrast to Shop A, Shop B had a higher number of purchases of smaller, 1-2 gram, less expensive items that were connected with gifts for new babies and family members. These did not need to be tried on in front of a mirror. We also learned that the owners of the second store generally did not like customers trying on a lot of jewelry (unless shoppers had bought their jewelry in the past) and encouraged customers to buy an item, wear it for a few days and exchange it if they did not wish to keep it. When asked, the owners did not think there was value in connecting customers' browsing behavior to their sales inventory, or to explore alternative ways of selling. Though their sales had dipped due to the recession, they were less willing than the manager in Shop A to explore alternative ways of making a sale. In fact, their sales people were more receptive to using CRM but were discouraged in using it by the owners. This may have been due to the fact that the owners in Shop B supervised the floor and weighed and priced jewelry items but did not personally sell to each customer, unlike the manager in Shop A who was intimately involved with each part of the sales process. To conclude, it was clear that a proportion of purchases need not involve people trying on jewelry, and that we needed to target stores where *browsing* activities were actively encouraged.

The aim of the fourth deployment was to test whether shoppers preferred to browse still or video images on the CRM. Originally, three South Asian couples and two other women who regularly bought jewelry were invited to use

the system in Shop B. At the last minute, the shop owners objected, saying that it would distract other shoppers. Rather than reschedule the session, we moved the prototype into a business acquaintance's home. Over the course of two hours, we video-recorded three women trying on multiple sets of jewelry, while the others gathered around them, readily manipulating the mirror for others, helping each other try on pieces, encouraging people to record alternate poses, and repeatedly dragging the slider bar backwards and forwards while they admired or joked about which earrings or chokers looked best.

Still vs. Video

Four of the five women said that they preferred to compare images of still photographs using the slider rather than playing back and pausing the video. The reason for this, as one person stated, was that she "didn't want to watch the whole movie" including the distraction of seeing themselves put on or take off their jewelry and unflattering shots of their arms or part of their face. They also wished to have a larger set of still images to compare against each other, rather than just two sessions at a time. By analyzing the video, we derived that the ability to manually use the slider to specify image frames for comparison provides random access to the useful segments of the image sequence and a lightweight way to avoid unwanted segments of the image sequence.

Serendipitously, privacy concerns were abated during the home deployment and participants felt freer to experiment with the system. Specifically, they recorded interactions for longer periods of time and lingered over the images and repeatedly visited multiple sessions. Several people said that they would prefer the mirror in the home where they could try on jewelry with different outfits.

Reaction to Image Matching Capability

Across all deployments, the image matching capability caused the most excitement among shoppers. Despite that the matches are inexact (sometimes dramatically) shoppers were impressed by the ability of the system to automatically match poses in images across sessions. The slider was the single most used widget in the UI, promoting a good deal of interaction among shoppers and their companions. The affordance of the slider bar on the touch screen was its high degree of interactivity to quickly make a comment about the appearance of an item in a particular image; this was true both for sales people and their customers. From these reactions, we concluded that the review component was indeed supporting interactions between shoppers and their companions and sellers. The image matching capability was viewed as a "cool" capability, and the excited reactions that we observed clearly demonstrated that the system contributed to the fun and engagement that it was intended to support.

A Third Person Perspective

An unexpected finding was that many participants noted that the captured images provided them with a greater sense of how a third person would perceive their look than they were able to gain from the mirror alone. Not only did the review component enable individuals to compare the pieces against each other, but more importantly to some respondents, they could envision how the item might appear from someone else's perspective, which is a perspective they did not obtain with a traditional mirror. The participants noted that there were times that they thought that a piece displayed on the wall would look good on them, and they maintained the same impression even after trying on the piece and looking at themselves in the mirror. However, a couple of people noted that the recorded images helped to dissuade them from buying something that did not look as nice as they had anticipated. Looking at recorded images of themselves provided a third person view which enabled shoppers to evaluate jewelry from a novel perspective.

This was surprising because the images that were recorded are nearly identical to what was seen in the mirror. We speculate that when using a mirror, the fact that the image in the mirror is of oneself is perhaps reinforced by the precise synchrony of motion in the mirror with one's body movement and thus when observing captured images, even ones where motion is matched between the images, it is easier to take on a third-person perspective. This effect warrants further exploration in future studies.

Recall

Respondents confirmed an expected affordance of the system: that it helped them recall what they had tried on, "It helps us remember what we wore." For example, when people tried on four or six items over the course of 20 minutes, they often had trouble recalling the appearance of the first few items. The system provided a convenient inventory of recordings well after the sales person had put items away. This validated that recording fittings is valuable to shoppers.

Salesperson and Shopper Interaction

One store manager found the prototype easy to use, and actively encouraged customers to try it. He would also use the time when he explained the system to his customer as an opportunity to make his sales pitch. One of the questions we had going in was the privacy issues of a camera recording images of shoppers. We found that because the system was deployed in the jewelry shop that customers would trust the store owner to properly maintain their privacy. This was an example of how the existing salesperson-shopper relationship helped to establish trust.

CONCLUSIONS

Our research aims are to create Ubiquitous Computing technologies that do not merely provide the same kinds of information available in online merchandising, but to understand the unique aspects of processes used in physical retail shopping and to supplement those with technologies that enhance those unique aspects.

The research reported in this paper is a combination of technological and social sciences with contributions that include, first, the paradigm of "matched access" as an extension to "capture and access," and second, the qualitative observations from multiple field trials. The observations uncover a few surprising effects that warrant further research, such as the perceived third-person perspective of the images that are nearly identical to images in the conventional mirror which is perceived as first-person perspective. There are other contributions in this work such as the characterization of South Asian retail jewelry practices and tools, along with an evaluation of image matching techniques that can provide "pose matching" for similar "matched access" systems.

The main objective in this paper has been to describe the complete process of examining an existing practice, identifying the information needs and creating a new technology to enhance the experience. Designing technologies that provide significant value to all parties in an application domain is a nontrivial challenge, particularly when the system must serve distinct and potentially conflicting user goals (seller and buyer). As Ubiquitous Computing technologies reach into new information-rich domains, such as physical retail shopping, designers should examine existing points of interaction among the parties and create novel technological capabilities that supplement and deepen the information used at those points of interaction.

ACKNOWLEDGEMENTS

We are grateful to our business partners for coordinating the in-store and home trials. We thank the store owners, employees, customers and all trial participants for feedback and inspiration. We also thank our colleagues at PARC for support and suggestions.

REFERENCES

1. ACEP Smart Mirror http://www.smart-mirror.com/en/produits.php?id=1 (accessed 8 Jan 2010)
2. Begole, B., Matsumoto, T., Zhang, W., Yee, N. Liu, J. and Chu, M., Designed to Fit: Challenges of Interaction Design for Clothes Fitting Room Technologies, *HCI International* (2009), in *LNCS 5613,* Springer, 448-457.
3. Bradski, G. and Kaehler, A. *Learning OpenCV: Computer Vision with the OpenCV Library.* O'Reilly Media, Inc., 2008.
4. Chaudhuri, A., Emotion and Reason in Consumer Behavior, Elsevier, 2006.
5. Commodity Online, India's Gems and Jewellery Market is Glittering. (20 Aug 2007) (accessed 8 Jan 2010) http://www.resourceinvestor.com/pebble.asp?relid=34936
6. Davis, F., *Fashion, Culture, and Identity*, U. of Chicago Press, 1992.
7. Danziger, P., *Why People Buy Things They Don't Need*, Dearborn Trade Publishing, 2004.
8. EZ Face virtual mirror http://www.ezface.com/EZface/Application/EZfaceApplication.aspx (accessed 8 Jan 2010)
9. Forsyth, D., Arikan O., Ikemoto, L., O'Brien, J., and Ramanan, D. Computational Studies of Human Motion: Part 1: Tracking and Motion. *Foundations and Trends in Computer Graphics and Vision*, *1*, 2/3 (2006).
10. Fujinami, K., Kawsar, F., Nakajima, T., "AwareMirror: A Personalized Display Using a Mirror", in *Proceedings of Pervasive 2005*, pp. 315-332.
11. Hawkins, D., Best, R. J., Mothersbaugh, D. L., *Consumer Behavior: Building Marketing Strategy*, 10th edition, McGraw-Hill Companies, (2007).
12. Iwabuchi, E., Nakagawa, M., Siio, I. Smart Makeup Mirror: Computer-Augmented Mirror to Aid Makeup Application. *HCI International* (2009) in *LNCS 5613,* Springer 2009, 495-503.
13. Lee, D., Park, J., Lee, M., Hahn, M., "Personalized Magic Mirror: Interactive Mirror Based on User Behavior" in *5th International Conference On Smart Homes and Health Telematics, ICOST 2007*, Nara, Japan, pp. 162-169.
14. Lindsay, G. Prada's High-Tech Misstep. *Business 2.0*, 1 (March 2004). (accessed 8 Jan 2010) http://money.cnn.com/magazines/business2/business2_archive/2004/03/01/363574/index.htm
15. Lurie, A., *The Language of Clothes*, Random House, NY, 1981.
16. McKinsey Global Institute. The 'B*ird of Gold'*: The Rise of India's Consumer Market. (May 2007).
17. Miller, D., *A Theory of Shopping*, Cornell University Press, 1998.
18. Miller, D., *Shopping, Place and Identity*, Oxon, Routledge, 1998.
19. Mistry, P., Maes, P., Chang, L. WUW - Wear Ur World - A Wearable Gestural Interface. In the *CHI EA '09 Extended Abstracts.*ACM Press, pp. 4111-4116.
20. Newcomb, E., Pashley, T., Stasko, J. Mobile Computing in the Retail Arena. *CHI 2003*, ACM Press, April 5–10, 2003, pp. 337–344.
21. Prus, R.C. Making Sales: Influence as Interpersonal Accomplishment. Sage Publishing, (1989).
22. Rodden, T., Rogers, Y., Halloran, J., and Taylor, I. Designing novel interactional workspaces to support face to face consultations. In *CHI 2003*. 57-64.
23. Truong, K., Abowd, G., Brotherton, J. Who, What, When, Where, How: Design Issues of Capture & Access Applications. In *Ubicomp 2001 Conference Proceedings*, ACM Press, (2001), 209-224.
24. Tran, Q., Mynatt, E. Cook's Collage: Two Exploratory Designs. In *CHI 2002*, (2002).
25. Underhill, P., *Why We Buy*, Simon & Shuster, 1999.
26. Xu, Yan, Mirjana Spasojevic, Jiang Gao, Matthias Jacob. Designing a Vision-based Mobile Interface for In-store Shopping. *NordiCHI 2008: Using Bridges*, 18-22 October, Lund, Sweden
27. Zhang, W., Matsumoto, T., Liu, J., Chu, M., and Begole, B. An intelligent fitting room using multi-camera perception. In *Proceedings of Intelligent User Interfaces (IUI '08)*. ACM, New York, (2008), 60-69.