Supporting Healthy Grocery Shopping via Mobile Augmented Reality

JUNHO AHN, JAMES WILLIAMSON, MIKE GARTRELL, RICHARD HAN, QIN LV, and SHIVAKANT MISHRA, University of Colorado Boulder

Augmented reality (AR) applications have recently become popular on modern smartphones. We explore the effectiveness of this mobile AR technology in the context of grocery shopping, in particular as a means to assist shoppers in making healthier decisions as they decide which grocery products to buy. We construct an AR-assisted mobile grocery-shopping application that makes real-time, customized recommendations of healthy products to users and also highlights products to avoid for various types of health concerns, such as allergies to milk or nut products, low-sodium or low-fat diets, and general caloric intake. We have implemented a prototype of this AR-assisted mobile grocery shopping application and evaluated its effectiveness in grocery store aisles. Our application's evaluation with typical grocery shoppers demonstrates that AR overlay tagging of products reduces the search time to find healthy food items, and that coloring the tags helps to improve the user's ability to quickly and easily identify recommended products, as well as products to avoid. We have evaluated our application's functionality by analyzing the data we collected from 15 in-person actual grocery-shopping subjects and 104 online application survey participants.

CCS Concepts: • Information systems → Information systems applications → Mobile information processing systems

Additional Key Words and Phrases: Augmented reality, mobile health, recommendation, grocery shopping, nutrition

ACM Reference Format:
Junho Ahn, James Williamson, Mike Gartrell, Richard Han, Qin Lv, and Shivakant Mishra. 2015. Supporting healthy grocery shopping via mobile augmented reality. ACM Trans. Multimedia Comput. Commun. Appl. 12, 1s, Article 16 (October 2015), 24 pages.
DOI: http://dx.doi.org/10.1145/2808207

1. INTRODUCTION

It has amply been noted that information technology can help catalyze a number of important benefits in healthcare that include improving its quality and reducing its cost [PCAST 2010]. With the emergence of sensor-rich, powerful smartphones that can provide a rich set of user contextual information in realtime, it has now become feasible to provide effective and affordable healthcare to nearly everyone via smartphones. In particular, carefully designed smartphone applications have the potential to enable individuals to participate in their care, which transforms healthcare systems from reactive to preventive, from clinic-centric to patient-centric, and from disease-centered to wellness-centered.

This article explores the use of smartphones, cloud computing, mobile augmented reality and related information technology to help improve societal health and wellness. Earlier research has shown a strong link between poor dietary choices and the increased risk of poor health conditions such as obesity as well as chronic diseases such as cardiovascular disease and diabetes. Poor diet and physical inactivity are the two most

Authors' address: Department of Computer Science, University of Colorado, Boulder, CO 80309-0430, Correspondence email: mishras@cs.colorado.edu.

DOI: http://dx.doi.org/10.1145/2808207

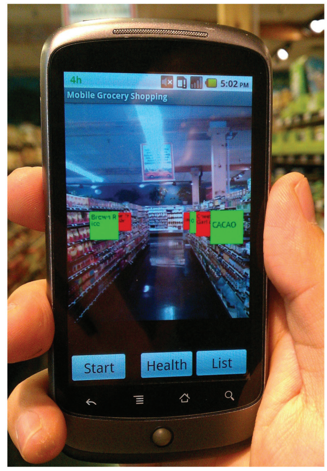

Fig. 1. A screenshot of our mobile application to assist in healthy grocery shopping. Augmented reality color tags identify healthy and unhealthy products.

important factors contributing to an epidemic of overweight people and obesity in the United States. Improving one's diet begins by improving the nutritional quality of the food choices he/she makes. In a food supply including tens of thousands of processed and packaged foods with diverse messaging on bags, boxes, bottles, jars and cans, making more nutritious choices is challenging at best for the average consumer [Katz et al. 2009]. Consumers claim to understand what is healthy and unhealthy, but acknowledge confusion over implementing general nutritional advice into practice [Lobstein and Davies 2009]. Providing consumers with nutrition information at the point-of-purchase has the potential to improve consumer decision-making about healthy foods, and thus have a greater impact on dietary quality than traditional generic messages of "eat better".

The use of technology in managing diets has been heralded as an effective tool and resource in helping to reduce the prevalence of poor health conditions and improve the general wellness of the public [Barton et al. 2006]. We propose to address the critical problem of improving the nutritional quality of the food choices individuals make by introducing mobile augmented reality (AR) at the point-of-purchase in grocery stores. AR is one of the most exciting emerging technologies, and in simple terms provides rich visual interaction with the real world by augmenting or overlaying a camera's view with computer-generated elements containing useful information relevant to the objects shown in the camera's video screen. With an AR-based smartphone application, a user can enjoy an instantaneous interactive or context-rich experience. AR has recently achieved significant mindshare as an exciting new technology for mobile smartphones. Examples include Golfscape GPS Rangefinder, an augmented reality range finder for golf lovers [Golfscape 2014]; DanKam, an AR application for people suffering from color-blindness [DanKam 2010]; Google Sky Map, an AR application for amateur astronomers [SkyMap 2011]; Word Lens which translates a foreign language captured by the mobile camera and overlays the result on top of the text [WordLens 2014]; and many more.

A prototype of our augmented reality mobile grocery shopping application is shown in Figure 1. As the user pans and walks up and down a grocery store aisle, the AR tags corresponding to highlighted products will change based on what products the user is facing. As a user walks towards an item along the aisle, its corresponding

AR tag grows in size. The tags when clicked reveal nutritional information about the product. The tags are also colored, for example, green to indicate products that are nutritionally preferable (e.g., low-calorie, gluten-free), and red to indicate products to avoid (e.g., high cholesterol or peanut content). Further, shoppers can specify health profiles which may impact their food purchase choices, such as weight control, heart disease, food allergies, etc. The recommended products shown via AR tags will then change depending on what health condition/concern is indicated by the user. We believe our system is the first to integrate augmented reality tagging and pedometry-based localization with a back-end server to provide health-based grocery recommendations at the point-of-purchase. We evaluated the effectiveness of our system in a real grocery store aisle with 15 actual grocery shopping subjects to determine how easy and fast the subjects reported it was to locate healthy food products and avoid unhealthy ones, using AR tagging with our application. We also evaluated our application's functionality and performance by analyzing data we collected from 104 online application demonstration/survey participants.

2. RELATED WORK

Augmented reality has been recently applied in the mobile health arena in a variety of applications. For example, AR tags are overlaid in a home environment to provide instructions to the elderly for tasks like taking medication, cooking, washing, etc. [Hervás et al. 2011]. TriggerHunter is an AR-based game that overlays tags on potential asthma triggers in the local environment [Hong et al. 2010]. Neither project contains any evaluation. An AR-based game has been developed for mobile phones to help individuals overcome insect phobias by allowing patients to kill virtual insects overlaid on the mobile screen [Botella et al. 2011]. A framework for several AR-based Q&A games has been created to rehabilitate patients [Lin et al. 2011]. An AR-based mobile game has been described that forces players to travel to various physical sites to obtain AR-overlaid information, thus facilitating exercise [Görgü et al. 2010].

Supermarkets are an excellent location to introduce informational [Ganapathy et al. 2011; Anderson et al. 2001; Wientt et al. 1997; Mhurchu et al. 2007] and dietary behavior [Kalnikaite et al. 2011; Mankoff et al. 2002] interventions because they are the place where most individuals in the United States make decisions and purchase their food products. An example of an informational intervention is a system where participants take pictures of items, for example, chips, which are then matched in an image database to provide product information that is overlaid on the picture of the product [Ganapathy et al. 2011]. This system requires shoppers to know exactly where the product is and still read the nutritional label on the packaging.

To aid individuals locate the items they are looking for and provide a high level health information about the products, visual guiding systems have been deployed in grocery stores and supermarkets. To direct the individuals to the items of their interest, these systems work in a hierarchical way, for example, a large sign of a general category such as "Produce" or "Dairy" visible from a distance followed by specific aisle signs about more specific item categories placed near the general category sign. Lately, these systems have started providing general health related information such as "Le Bio" or "Le Bonne" in Carrefour stores or "Organic" in Safeway stores. While these visible guiding systems certainly help individuals in making healthy choices, a key limitation is that they provide generic information and are not tailored for each individual based on his/her health condition and other factors. We compare the performance of our AR-assisted mobile grocery app with a visual guiding system in a real grocery store in Section 5.4.

Other informational interventions rely on shoppers stopping by a supermarket kiosk to receive nutritional information [Anderson et al. 2001; Wientt et al. 1997] and coupons to incentivize healthier choices [Mhurchu et al. 2007]. Although these systems did encourage participants to purchase healthier food, marginalized populations were less likely to use the system. Ubiquitous grocery intervention systems are promising for dietary behavior change because they are always with the shopper and can provide just-in-time information about food items. For example, Mankoff et al. [2002] designed a system where shoppers could scan their receipts and receive information about the nutrition of the items. This system provides shoppers the ability to reflect on the food that makes-up their diet after purchasing the foods. Other point-of-decision ubiquitous computing applications for grocery shopping describe ways to use LEDs to inform the user either about nutrition through a small clip on stick [Mulrooney et al. 2006] or about how many miles the food traveled via a device clipped onto grocery carts [Kalnikaite et al. 2011]. Similar to other work discussed, these systems require the user to know the location of the item and to select it to gain information about the item. Our application makes real-time customized recommendations of healthy food items to get and unhealthy (or harmful) food items to avoid, and AR-assisted color tags to facilitate healthy food purchase decisions.

Recommender systems have been an area of active research for decades and many techniques have been proposed (see Bobadilla et al. [2013] for a survey). A number of food recommendation techniques have also been proposed recently, such as recipe recommendation [Freyne and Berkovsky 2010], context-aware food recommendation at table [Oh et al. 2010], and food recommendation for people with diabetes [Phanich et al. 2010] or tourists with certain health concerns [Agapito et al. 2014]. In this work, we aim to recommend/warn shoppers of grocery items in the current isle based on personal and family health profiles and grocery items' nutritional information.

Pedometry-based navigation using accelerometer data from mobile phones provides a convenient and low-cost way to monitor user progress up and down a grocery aisle without requiring an extensive localization infrastructure. A variety of step estimation algorithms have been proposed [Fuchs et al. 2011; Ladstaetter et al. 2010; Beauregard 2007]. For our purpose of aisle navigation, we adapt a simple approach that achieves sufficient accuracy using personalized pedometry by estimating individual stride length [Ahn and Han 2011].

3. APPLICATION DESIGN OVERVIEW

Our goal is to build an indoor mobile augmented reality system for healthy grocery shopping by leveraging the sensing and AR capabilities of smartphones and the knowledge of health rules in order to recommend appropriate products to purchase or identify products to avoid. We seek to understand two basic questions.

—How much time does AR tagging of recommended products save a grocery shopper with a given health condition in comparison to the current approach of preparing a grocery shopping list?
—Does highlighting unhealthy products help the user to reduce the time it takes to confirm avoidance of items that would conflict with their health condition?

3.1. Design Requirements and Assumptions

Our system needs to be able to support navigation within a grocery store aisle. It needs to provide AR-based tags that are geographically (i.e., shelf location along an aisle) associated with recommended products, or products to avoid. The recommendation of healthy or unhealthy products needs to be determined in real time. The system should measurably improve the shopping experience of the health-motivated shopper, whether measured by the reduction in time to find their desired products, or by an improved ability to avoid unhealthy products. The system should be relatively easy to

Fig. 2. The coordinate system we use for item locations in a grocery store aisle.

use and learn. Also, the system should leverage existing low-cost sensors on most mobile devices, and not require a costly in-store infrastructure, such as an infrastructure for localization.

Based on our discussions with local grocery stores, we find that grocery stores have an electronic product database, though not necessarily associated with location. Note that we do not require exact location of items on store shelves, but only approximate sectional information along an aisle, quantized as shown in Figure 2. Given such a coordinate system, we can overlay tags of healthy/unhealthy items coarsely by section, which should be sufficient for the user to find the items quickly. We demonstrate that this coarse quantization is sufficient in our evaluation.

To better understand users' grocery shopping behavior with respect to their food product purchases, we conducted an online food-shopping demo and survey of 104 human subjects. This research was approved by Institutional Review Board (IRB) [Ahn et al. 2012]. The demo and survey consisted of three steps that participants were required to complete.

(1) Participants find four healthy products of their choice, that had low calorie and no milk content, from among 60 picture-based grocery products displayed on the website.
(2) Participants view a 3-minute video demo to familiarize themselves with our shopping app.
(3) Participants provide feedback, including an evaluation of the brief picture-based shopping experience, an evaluation of the video demo, and a detailed feedback on their personal grocery shopping behavior, focusing specifically on healthy food shopping.

We designed our online survey, using GoogleDocs' online survey tool [GoogleDocs 2014] and deployed it on Amazon's public MechanicalTurk website [MechanicalTurk 2014].

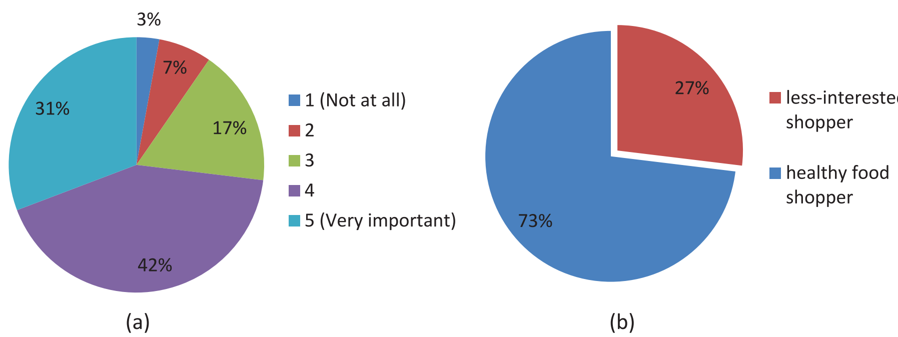

Fig. 3. Result of importance-rating for buying healthy food products (a) on a five-point scale, and (b) on a two-point scale from healthy food shopper to less-interested shopper.

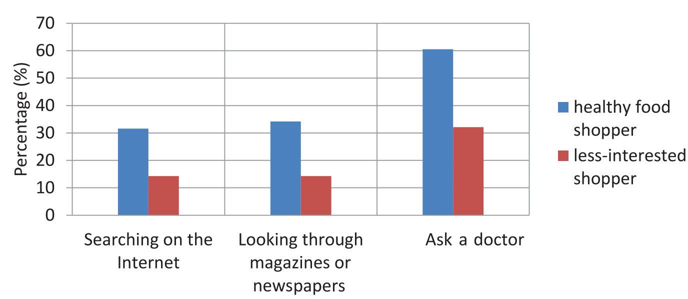

Fig. 4. Comparison of the two groups' pre-grocery-store visit healthy food searching behaviors.

3.2. Users' Grocery Shopping Behavior

3.2.1. Healthy Food Shopping Behavior. We categorized participants into two groups: those who are more interested in buying healthy food products and those who are less interested in buying such products. The reason to categorize people into these two groups is because people who are interested in buying healthy food products are likely to be much more focused on healthy eating and ensuring that they can obtain healthy food products frequently and easily. Our goal in making this distinction between these two groups was to see if their shopping behaviors were distinct from each other and how far apart ratings of their shopping behavior patterns would actually be. These results can therefore help us to better understand all users' grocery shopping behaviors and help us to further evaluate and improve the use of our application. We asked the following question: How important is it to you that you buy healthy products (e.g., low calorie, low sugar, organic, etc.) for yourself and/or your family when you go grocery shopping?

We categorized the two groups as follows: healthy food shoppers – those who provided ratings within the 4- to 5-point range ($n = 76, 73\%$); and less interested shoppers – those who provided ratings in the 1- to 3-point range ($n = 28, 27\%$) as shown in Figure 3(b). We investigated three different food grocery shopping behaviors for members of these two groups: pre-grocery store visit searching behaviors for healthy food products; preferred methods searching for healthy products in a grocery store; and food quality factors considered most important when choosing healthy products. Figures 4, 5, and 6 show a comparison between the two groups for these grocery shopping behaviors.

Our first finding is that the healthy food shoppers spend almost twice as much time as the less-interested shoppers in using different search methods to search for healthy products prior to their grocery store visits. "Asking a doctor or friends" was

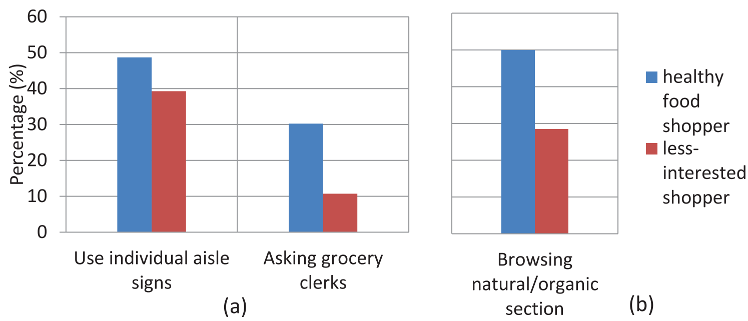

Fig. 5. Comparison of the two groups' in-grocery store healthy food searching behaviors (a) for finding unknown location products and (b) for finding healthy food products.

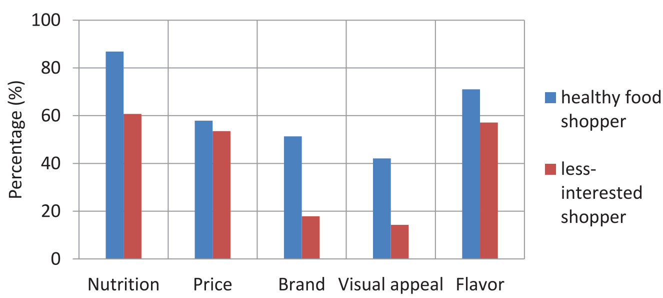

Fig. 6. Comparison of the two groups' healthy food quality preferences in the picture-based demo.

the most frequently noted method of preference for all users in gathering information on healthy food products. Our second finding is that when subjects need to find a grocery product of interest, healthy food shoppers more frequently prefer using the aisle signs or asking grocery clerks than do the less-interested shoppers. Also a much larger proportion – three times as many healthy food shoppers as compared to the less-interested shoppers – preferred asking grocery clerks directly for help, when trying to locate food products of interest. Additionally, when the subjects needed to locate a healthy food product, almost twice as many healthy food shoppers as the less-interested shoppers, chose to browse the organic/natural food sections of the grocery store. Finally, our third finding is that, while the order of importance for each of the food quality factors was the same for the two groups, "nutrition > flavor > price > brand name > visual appeal", a much larger percentage of healthy food shoppers considered nutrition, brand, and visual appearance as highly important food qualities in selecting a health food product than did the less-interested shoppers.

3.2.2. Nutrition-Based Multiple-Choice Data Collection. Our first approach to investigating the kinds of healthy food content information our application should provide involved collecting online and in-person survey information from our initial project 25 participants as to the types of health conditions, diseases or food sensitivities, they or their family members need to consider when shopping for food products. The 25 participants and/or their family members had at least one health problem and 16% ($n = 4$) of them had more than two health problems. We asked the participants about their own health conditions and their family members' health needs, since we assumed that many shoppers often shop not only for themselves, but for other family or household members as well. We found that 79% ($n = 82$) of the survey subjects usually buy grocery food products for their family members when they go shopping. These two cases showed that our mobile application would need to provide in-depth information on food products that were suitable for a large variety of multiple diseases and food sensitivities all at one time to the user, while she/he was shopping.

Since, there are just too many possible health conditions and combinations in the real world, providing correct recommendations for each and every possible heath condition is impractical. So, we refocused our approach to a solely nutrition-based approach. In this way, regardless of the specific health condition or conditions of users, our application would be set up to query users about the nutrition content of food products they need or wish to purchase. For instance, most users with a known specific health condition, have already been advised by their doctors as to which food products to avoid or select that will ameliorate their condition – for example, a person with diabetes would already have been advised to eat food products with low-sugar content. Later in our application's development, in surveying potential users of our application, we found this approach to be corroborated by their feedback, when the largest percentages of both the healthy food shoppers (87%) and the less-interested shoppers (60%) indicated that nutrition was the highest rated food quality factor of interest to them when purchasing grocery food products, as shown in Figure 6.

3.2.3. Preference of Aisle-Based Display. From the online survey data, we also found that potential users of our application prefer an aisle-based AR display of grocery food products on the smartphone over other types of displays. Participants were asked to rate their preferences for three different AR displays of healthy/nutritional food products recommended by the mobile application on the smartphone: all grocery store products at once, one aisle's products only, one section of an aisle's products. Over half ($n = 59, 57\%$) of the participants indicated preference for displaying recommended food products in one aisle only. Also, the online survey subjects indicated that they have frequently bought additional food products they were not originally planning to buy, which were located near the product they were buying. This result indicates that another benefit of an aisle-based AR application display is for users to be able to evaluate more quickly nonplanned food purchases in the grocery store.

3.3. System Overview

Our system consists of an external image labeling service, a mobile cmponent, and a remote cloud server component (see Figure 7). To determine the initial location of a user in the grocery store, the mobile component sends a product snapshot to the cloud server component, which forwards that snapshot to an external image labeling service. This external image labeling service returns the product identity to the cloud server, which then determines the current location (aisle) of the user by referring to an indoor layout of the grocery store. After determining the identity of the aisle in the grocery store, the mobile component estimates user motion, thus providing a position estimate within the aisle, as well as orientation. The user also inputs his/her health profiles on the mobile client, for example, seeking some combination of low-calorie, low-sodium, low-fat, lactose-free, nut-free, etc. items. This position estimate along with orientation and health condition is then again communicated to the server, which consults the product location database along with its health rules to come up with a recommendation of products to buy or avoid. The server has access to the nutrition facts and ingredient lists of products, and can thus apply health rules to decide whether products are healthy or not for the given health condition(s). These highlighted items are then sent to the mobile client, which renders the recommendation results on the screen via AR. In order to achieve real-time performance, it is helpful for the mobile application to cache data items locally on the client, so as to avoid excessive network communication latency.

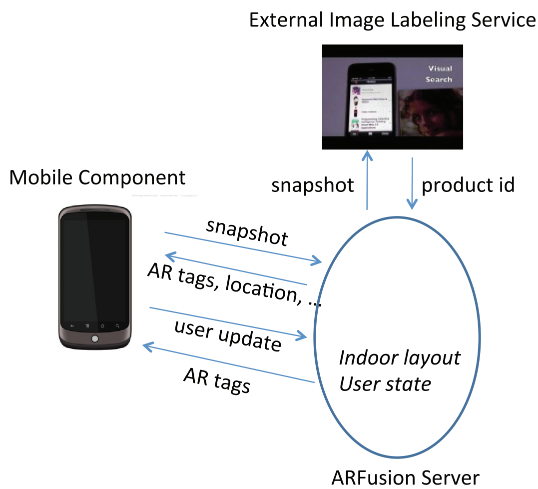

Fig. 7. System architecture.

In order to meet the objectives of low-cost, low-infrastructure navigation within a grocery store aisle, we use three-axis accelerometer information obtained from a user's mobile phone to estimate the distance traveled by a walking user. We build one such personalized pedometry system into our application.

In order to match the 3D perspective of a supermarket aisle and the intuition of the user about how far away a tagged item is located, the AR tags for items that are closer to the user are rendered larger than the tags corresponding to more distant items. The result is that as a user navigates down an aisle, tags grow in size until they pass by out of view as the user walks past the item, thus giving the user a 3D AR experience.

In order to clearly differentiate between healthy and unhealthy products in the user interface, we have used intuitively colored tags: green for good/healthy products; red for products to avoid. We measure the effectiveness of this approach in the evaluation section. Additional mappings of colors to different categories of food were considered, for example, vegetables, meats, dairy, fruits. We found that the latter approach was confusing, therefore we focus only on the color tagging of healthy/unhealthy food products.

4. SYSTEM COMPONENTS

4.1. Image-Based Positioning

Our system requires accurate determination of the user's location in an indoor environment. Locating the user in an indoor environment using the hardware available on a smartphone is a challenging problem. The Global Positioning System (GPS) cannot be used in indoor environments, since line-of-sight communication between GPS receivers and satellites is not possible in an indoor environment. Radio frequency (RF) positioning systems that use WiFi and Bluetooth radios on smartphones provide limited accuracy (1 - 3 m) due to the complexities associated with indoor environments, including a variety of obstacles (people, furniture, equipment, etc.) and sources of interference and noise from other devices [Gu et al. 2009]. Some of these RF positioning systems use RF location fingerprinting, which requires relatively time consuming site survey that may not be feasible for large indoor shopping environments. Therefore, we investigated the use of other positioning technology.

Our system uses a commercial image labeling Web service, called IQEngines [IQEngines 2013], to determine the user's initial starting position when using our grocery shopping application. IQEngines uses a combination of computer vision and crowdsourcing to tag a photo with a label describing the content of the image. For example, an image of a box of Frosted Cheerios cereal might be labeled "General Mills Frosted Cheerios". When an image is submitted to IQEngines, the image is first processed by a computer vision system in an effort to provide an accurate label. If the computer vision system cannot identify the image, then IQEngines passes the image to its crowdsourcing network for analysis and tagging. According to IQEngines, the time to return a label for an image varies from a few seconds for the computer vision system, to a few minutes for the crowdsourcing system. To ensure fast image labeling in our experiments, we have pretrained IQEngines with specified images and associated labels for each of the food items in our test environment.

To locate a user within the indoor shopping environment, our mobile application prompts the user to take a picture of the nearest food item using the smartphone. After this image is submitted to our cloud server, the server submits the image to IQEngines for labeling. Upon receiving the item label for the image, our server looks up the location for this item using a spatial database. This spatial database contains the name, location, and other associated metadata for each item found in the shopping environment. In our grocery shopping application, the coordinate system for item locations is expressed using the following dimensions: aisle number, aisle side (left or right), division number, shelf number, and item sequence number. Based on our conversations with local supermarkets, we have found that this coordinate system is representative of item databases found at some establishments. In this coordinate system, aisles are separated into 4-foot divisions, and shelves in each aisle are numbered from bottom to top. Items in each location specified by a tuple of "aisle number, aisle side, division number, shelf number" are ordered according to item sequence number. Figure 2 shows a graphical representation of this coordinate system for a typical grocery store aisle.

4.2. Localizing the User within the Grocery Aisle

As mentioned earlier, our approach is to apply prior work in personalized pedometry to the aisle navigation problem. However, some adaptation is needed for the grocery store scenario. We observed shoppers' behavioral patterns as they used our application in grocery stores. We found that the users did not always hold the phone consistently, upright, pointed forward down the aisle. The mobile phone's orientation was often changed, whenever they moved towards food items recommended by the AR tags of our application. When they moved towards the products they wished to purchase, they usually changed the mobile phone's orientation, such as holding it down by their side while walking or in a strange angle while holding a basket or operating a cart. Whenever this happened, the accelerometer sensor incorrectly detected a stride. To avoid these false strides, we modified the pedometry algorithm to ignore sudden changes in acceleration if they also corresponded with sudden changes in the orientation sensor.

Another modification we made to the pedometry algorithm was to limit motion to the component in the direction parallel to the long axis of the aisle. This is essentially 1D map matching, wherein the walls of the aisle form a map that confines the travels of the user to a set of acceptable paths, or in this case a single path. In this way, our algorithm cannot misestimate the user as being located within a shelf/wall, and thus our location error is limited to lie only along the long axis of the aisle.

To achieve this, we construct a bounding box around each aisle, where we bound the range of the x-axis by the width of a regular aisle in the grocery. When the user approaches the edge of the bounding box, for example, the shelves, then we only take the component of the motion along the axis parallel to the bounding edge, and ignore

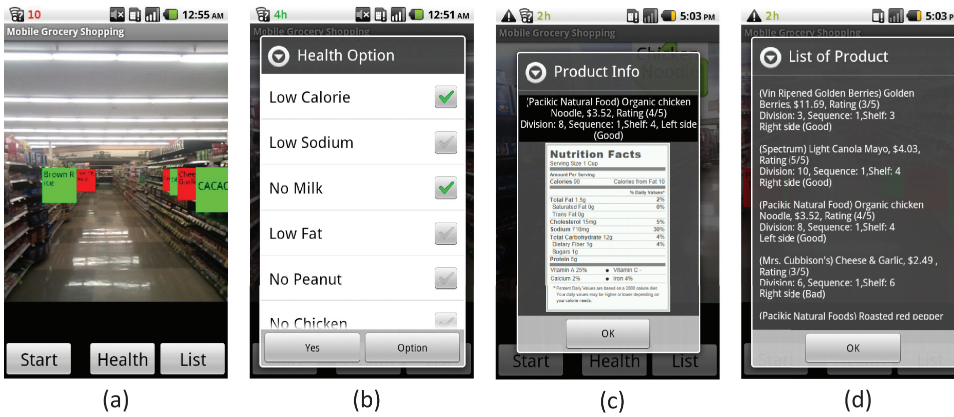

Fig. 8. Screenshot of (a) the AR mobile shopping app, (b) the health conditions selection screen activated by clicking the "Health" button, (c) the product information screen activated by clicking on an AR tag associated with a product, and (d) a typical non-AR grocery list used to compare against the AR UI (see evaluation in 5.4).

any component of motion perpendicular to the bounding edge. This approach keeps the user inside the bounding box. In this way, we were able to substantially improve the accuracy of our pedometry-based localization.

4.3. AR-Based User Interface

Our AR-based user interface is shown in Figure 8(a). AR tags are shown in 3D depth perspective, and are rendered using the OpenGL library. Products that are closer to the user will have larger tags, while products that are farther away will have smaller tags. To localize the tag next to the related product, we compared the distance on the phone between the product and the user with a distance on a real setting. The depth perspective was adjusted accordingly. Since the tags in our application are displayed in 3D space, we are able to adjust the display of the tags according to the angle at which the user is viewing an item using the phone. When the user looks at the front of the aisle, the tags are shown facing the user. If the user turns to the left or right to inspect a particular part of the aisle, the tags are automatically rotated to face the user.

In terms of hardware requirements, we found that a phone such as the Sony Nexus One, which has a 1-GHz processor, 512-MB memory, and 4-GB disk, was sufficient to run the OpenGL library to render AR objects in real time. In comparison, we found that running our application on an older Android phone, a TMobile Mytouch 3G running at 512 MHz, resulting in jerky rendering of AR objects, even after we upgraded from Android 1.6 to 2.2.

4.3.1. Dietary Food Constraints. People who have diabetes, allergies, hypertension, or other such health issues must often carefully monitor the types of food they buy in the grocery store. For instance, people with diabetes need to control their sugar level, so they must avoid high-carbohydrate food products that are high in sugar. People who have allergies, such as peanut or milk allergies, must purchase food products that do not contain these specific ingredients. People with hypertension should always try to avoid high-sodium products, in order to maintain their health. Some people may have multiple diseases or health issues (e.g., diabetes and allergies). Our application can help people or patients, under the care of a doctor, to monitor their food purchases and intake according to specific health issues they have. Figure 8(b) shows how a mobile user, with this application on their phone, can select different food ingredient requirements – such as "low calorie", "low sodium", "no milk", "low fat", etc.– that are tailored to their specific dietary needs. For example, a user with diabetes, hypertension, and milk allergies would choose “low sugar”, “low sodium”, and “no milk” on this screen. The mobile application then displays the actual food products on the grocery store aisle that are advised or unadvised to buy, with green and red AR tags, respectively. The application saves these settings, so the user only has to enter them on first use, though the user can change these settings any time. Also, the application supports multiple conditions, that is, if two conditions are checked and must be avoided, then all recommended products must satisfy both conditions.

As the user walks down the grocery aisle searching for products, she or he can easily get more information about an advised product or a product to avoid by tapping on the AR tag corresponding to the product that is displayed on the mobile phone. The information displayed when the user taps on the tag includes the product’s brand name and brief description, the nutritional information (FDA info), the price, location information (shelf number), and the selling rating – related to the store’s record of the frequency of purchase for the item. Indirectly this comprehensive information about the product’s content also provides an indication of the food product’s known or expected flavor. The condition of the actual product (e.g., fresh or wilted vegetables) and the manner in which it has been displayed on the shelf in the grocery store contributes to the user’s impression of the product’s visual appeal. These food quality factors and ingredients were identified as very important to the survey subjects who evaluated our application as potential real-life users of our system. The graph in Figure 6 shows the different food quality factor ratings that the survey subjects gave for their evaluation of factors they pay most attention to when selecting healthy food products. Figure 8(c) shows an example of the nutritional information displayed when the user taps on the mobile application’s AR tag.

4.3.2. Static- and Dynamic-Motion AR Tag Display. The ARFusion application provides users readable information on the phone regardless of their walking states. When a user walks down the grocery aisle with the mobile phone looking for preferred products, the tags on the phone would normally be shaking from the motion of the user. The user can have difficulty reading the information displayed when she taps on a tag. To correct for this, we propose two features, static- and dynamic-motion tag display. First, the feature for the static-motion tag display is used when the user walks down the aisle. The application displays the tags in fixed positions whenever the user points the phone in front of him in the same direction of the aisle the user is walking down. It provides a static display to the user that the tags can easily be read at fixed positions on the screen, even though the background may be varying wildly. Second, the feature for dynamic-motion tag display is used when the user is standing approximately stationary on the aisle. At this point, when the user pans the screen and points the camera at a product on a shelf, the screen allows tags to change position on the screen and rotate properly to face the user. To implement this policy, we checked the accelerometer every second to detect if there is motion or not, and adapted the AR display accordingly.

4.4. Health-Based Grocery Recommendation

Our AR-assisted mobile grocery shopping application is designed to make customized recommendations of healthy grocery products to end users in real time. The recommendations need to be customized since shoppers may have different health concerns such as food allergies, heart disease, or weight control. The recommendations also need to be generated in real time (while shopper is in a specific aisle) for them to be useful. The primary components involved in the recommendation process include the product database, shoppers’ health profiles, and recommendation strategies.

Products Database. This database maintains a variety of information regarding each product item in the grocery store that may be considered for recommendation. This database is usually populated by the store, but extra information may be obtained from manufacturer or online databases. Specific information of importance includes product name, ingredients, nutrients, as well as its location in a particular aisle (e.g., shelf section, level). Since the product items differ significantly in terms of ingredients and nutrients, we only consider the ingredients that people may be allergic to and categorize nutrients into coarser but more intuitive categories such as low calorie, low sodium, etc.

Health Profile. In order to recommend certain items to a shopper, the system must understand which items are required or wanted by the user. Furthermore, the system is capable of advising the user against the selection of certain items that have certain nutritional qualities or contain ingredients that may be harmful to him or her (e.g., ingredients to which the user is allergic). A simple health-based nutrition model was implemented to support these functionalities for testing the system. The model was populated with data from two main sources: (1) personal health-based profiles of users, for example, food information and ingredients that a person concerned with his/her weight and who also has a milk allergy might want to purchase for his/her diet or avoid altogether; (2) family health-based profile, for example, food qualities (e.g., calories, fat content) family members might prefer and ingredients that the family members may have been advised by doctors to avoid.

Recommendation Strategy. Food recommendation in grocery shopping environments is essentially a "matching" process between a shopper's health profile and certain food items in the products database. Based on existing dietary guidelines (e.g., DietGuideline [2010]), we construct a number of matching rules targeting different health profile categories and the corresponding food categories to recommend or avoid. At runtime, given the shopper's health profile and current aisle location, the server constructs a list of food items, each with one of two recommendation labels: *recommended* means the item has nutrition needed by the shopper, and *warn* means the item is in the list of harmful foods associated with the shopper's health profile. The recommendation results are then delivered to the shopper's mobile device for rendering. Note that our recommendation focuses on satisfying the rules based on the dietary guidelines. While not a focus of this work, more detailed recommendation strategies can be developed to consider other factors such as food price, taste, brand name, etc.

5. EXPERIMENTAL RESULTS

5.1. In-Person Survey Design

In order to validate our system, we collected in-person feedback from 15 users. These users provided us feedback in a couple of ways: First, they were asked to take an online survey so we could collect some basic demographics and information about their shopping needs and habits and any specific health/dietary restrictions. Second, users did an in-person survey with the researcher, after having accompanied the researcher while shopping in a grocery store for one hour and using our system on an Android phone. The grocery store we used for our experiments is Lucky's Market located in Boulder, CO. In this way, we were able to receive immediate verbal feedback from the subject on how easy and useful our system was to operate. Finally, the users completed a satisfaction survey, evaluating how the use of our system could potentially meet their needs for an improved healthy shopping experience. This user study was approved by the Institutional Review Board (IRB) [Ahn et al. 2012].

5.2. Pedometry-Based Localization

The pedestrian localization via pedometer and heading estimation systems were implemented and tested in Java on a Nexus One smartphone running the Android 2.2 (Froyo)

Table I. Step Detection Accuracy

Stride Length	Measured Steps	Actual Steps	Error (%)
Short 1	28	30	−6.67
Short 2	29	30	−3.33
Short 3	28	30	−6.67
Regular 1	30	30	0.00
Regular 2	30	30	0.00
Regular 3	29	30	−3.33
Long 1	31	30	3.33
Long 2	30	30	0.00
Long 3	32	30	6.67
$\|Avg.\|$			3.33
Std. Dev.			4.41

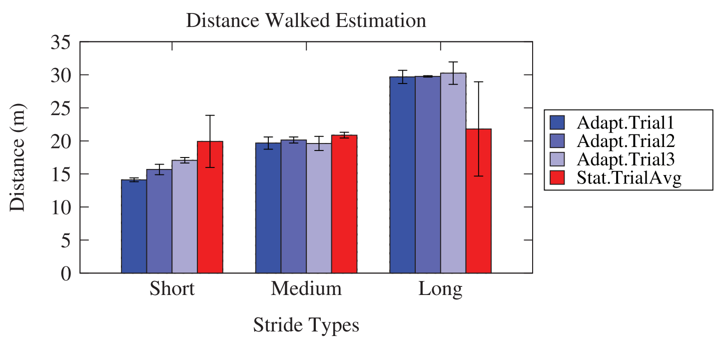

Fig. 9. Overall distance walked.

operating system. User tests to evaluate pedometry step detection, stride estimation, and the combination of step detection and stride length estimation into an overall distance walked estimation were carried out. Additionally, different types of users were simulated, from an "engaged" user who wishes to learn how to use the system to obtain the best performance, to the "casual" user who is not interested in performance and so uses the system in a careless manner. Further, the method to identify grocery-aisle angle using linear regression on user location history is evaluated.

To evaluate pedometry-system step, stride, and distance accuracy, a user was tasked to walk three trials of 30 paces in testing each of three different types of user strides. The first stride type is a "short stride," which is a deliberately short stride of about 50-55 cm. The second stride type is a "medium stride," which is a comfortable stride length of about 65-70 cm, which is natural for most users. Lastly, the "long stride" is one that is the largest the user can manage without jogging or running; a length of about 95-100 cm. The results of these nine trials can collectively be seen through Table I and Figures 9 and 10. Table I shows our system to have an overall step-detection error rate of 3.33 percent. In fact, for longer tests that we omit here, step-detection accuracy was shown to improve with the number of strides taken.

In Table I, short strides have tendency for under-detection, while long strides are prone for over-detection. This is due to the static threshold used for detection, which is tuned for the normal stride length scenario. An adaptive step threshold detection scheme was implemented and tested, but suffered a poorer performance than the static method. We theorize this counterintuitive result to be due to the accelerometer's 10-Hz maximum sampling rate on the Nexus One smartphone not providing a smooth enough data curve for the adaptive algorithm to leverage effectively.

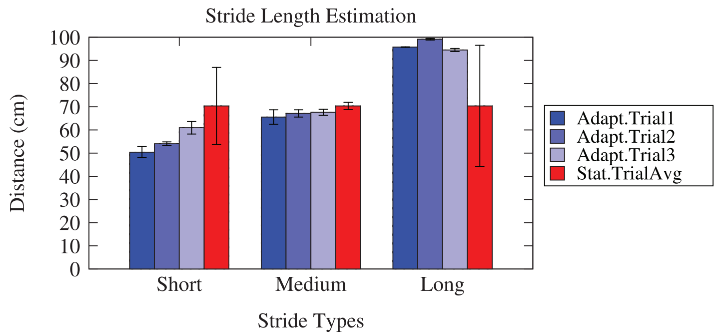

Fig. 10. Stride length estimation.

Figure 10 compares the static and adaptive stride length estimation techniques. The resulting stride lengths represent the average stride length of each of the nine user trials completed, calculated by the overall distance measured divided by number of steps detected, but not actually taken. This removes any additional step-detection errors that might be present and allow a pure comparison of stride length estimation. Not shown in this figure is the adaptive stride estimation overall error rate of 2.33 percent, while the static stride estimation suffers 17.06-percent error. Interestingly, because the static method was tuned for the medium stride length, its average error actually outperforms that of the adaptive method on the same data set. A point of note is the extreme accuracy of the long stride under the adaptive estimation scheme. The error bars are almost too small to be seen, averaging to 99.6-percent stride-length accuracy for this stride type. This excellent accuracy is most likely due to the flatness of the alpha correction function for large positive peak amplitudes.

Figure 9 addresses the combination of error from step detection as well as stride estimation techniques. An overall walk distance is measured by our system and is compared against the ground truth walked distances. In some cases, for example, adaptive trial 1 for a short stride, an error in step works to reduce the error stride. However, in most cases, if both kinds of error are present, they combine with one another, which is evident by the increase in overall error from stride (2.33 percent) and step (3.33 percent) to distance walked (3.43 percent).

After evaluation of the general step, stride, and distance performance characteristics of our system, we turn our attention to the operation of the system given the constraints of our target environment – the indoors of a grocery store. We look into the challenges of keeping user location in the aisle of interest, given complex user movements, and learning about the structure of the indoor environment – for example, grocery-aisle long-axis orientation – given only the motion sensors of the smartphone. To our knowledge, no other smartphone pedometry system has addressed the challenges of irregular and highly dynamic movement scenarios capable when a user is browsing during shopping. To this end, we evaluate three representative scenarios of possible user movement patterns that vary widely in possible user movement type and complexity. In doing so, we additionally stress-test our pedometry-bounded location method. Further, we examine the benefit of our bounded method in finding the long axis of a grocery store aisle, which can be done without any knowledge of the unique floor plan of the particular store our user is visiting.

We explore three representative scenarios: a 35-m walk down the long axis of a grocery store aisle with-return (casual walk), a repetetive circular walk of 2.5-m radius (circle walk), and a bowtie-shaped walk simulating a psuedorandom walk (bowtie walk). This last scenario, shown in Figure 15 and 16, also serves a second purpose in that we

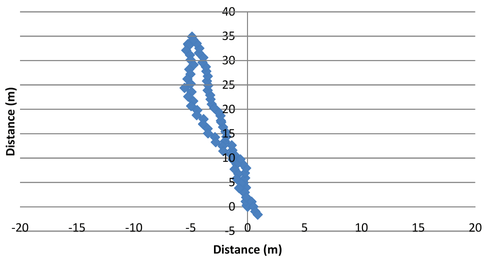

Fig. 11. Casual walk: Unbounded method.

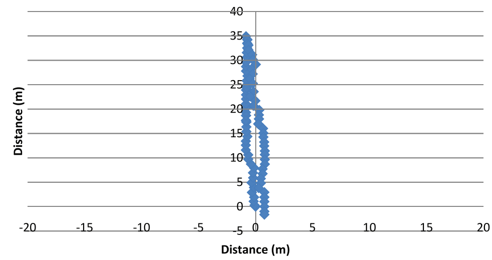

Fig. 12. Casual walk: Bounded method.

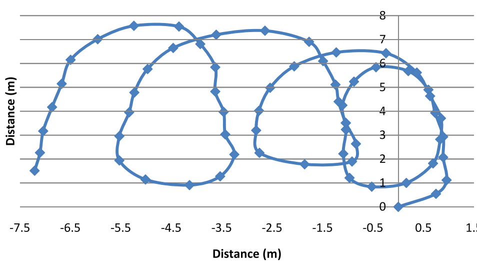

Fig. 13. Circle walk: Unbounded method.

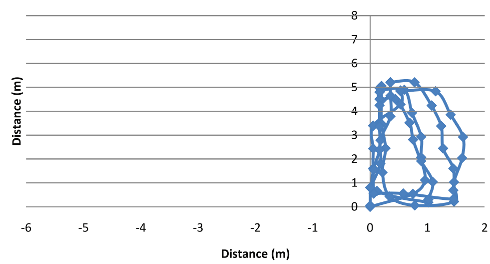

Fig. 14. Circle walk: Bounded method.

additionally use the trace to simulate a user's behavior prior to entering an aisle for the purposes of shopping under the use of our application – so that we may test our method of finding the aisle long axis orientation information.

First, we explore the effect of a "casual" user on a walk down and back the length of a 35 m long mock grocery store aisle. The test was carried out in lengthy hallway, and so this allowed us to stress-test the system by using a distance longer than is actually found in normal grocery-store aisles. The user type tested is classified as "casual," because for this user type, care is not taken to hold the smartphone in a verticle orientation, which would offer the highest locationing accuracy. Instead, this user is allowed swing the arm holding the smartphone, introducing a high level of noise data to the sensed user motion. An "engaged" user type was also tested in this scenario, but it is interesting to note that because the engaged user takes care in obtaining the best performance from the system, the bounded method was completely unnecessary in offering correction to the location information.

It can be seen in Figure 11 that the user drifts. This drift is caused by both the casual nature of the user type, as well as an inaccurate estimation of the grocery-aisle angle, that is, the aisle orientation was set to the left of true by 10 degrees. The bounded model, Figure 12, shows our systems corrective action under such a scenario. The total distance walked is shown to be shorter under the bounded method, so some sacrifice in accuracy is shown to be incurred, however such increased orientation accuracy is an acceptable tradeoff for the distance penalty.

Figures 13 and 14 test our system for an erratic circular user walking pattern. Often, a browsing shopper will return to a location of interest after initially passing it by. This scenario is very difficult to handle as errors in orientation cumulatively add at each step. In our test case, we use an engaged user, walking in a circle for four laps. The unbounded method in Figure 13 shows the effects of such orientation drift in which the user's virtual location would move across aisle boundaries. The bounded method handles this scenario very well. By forcing the user's location to be confined to a specific block of floor space, we avoid such drift.

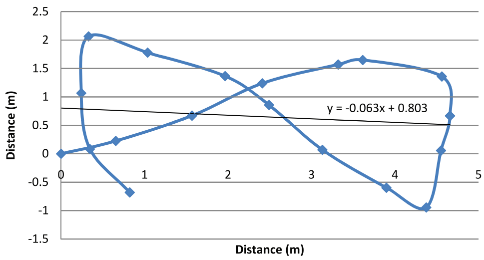

Fig. 15. Bowtie walk: Unbounded method.

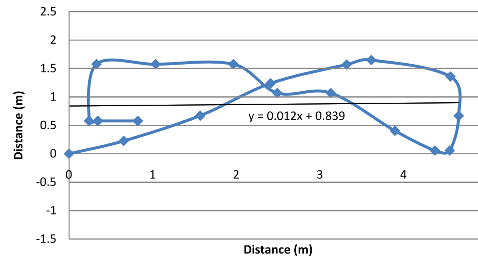

Fig. 16. Bowtie walk: Bounded method.

Finally, Figure 15 and 16 simulate a psuedo-random walk, useful for testing cases in that a user changes orientation direction more than once, which is common under a browsing movement. Also useful is this case to simulate user motion before entering an aisle. In Figure 15, we see the true axis of movement is off from the smartphone estimation – if not, the bowtie shape would not be tilted slightly to the right. The bounded method in Figure 16 again corrects for this, while incurring small finer-grained error as a tradeoff. Further, we simulate the aisle orientation algorithm by plotting a linear regression of step points overlayed on each bowtie shape. We show with a minimal collection of points our aisle long axis orientation estimation has good accuracy as shown by the line regression overlayed in the figures, therefore providing higher accuracy for use in the system's postimage localization mobile AR shopping phase.

5.3. Image-Based Positioning

The accuracy of the IQEngines service was tested through taking pictures according to varying angles, as shown in Figure 17. Ten grocery items were photographed at 45, 0, and -45 degrees and the pictures were then sent for evaluation to IQEngines, which then reported back its result. We took one picture straight from the front of the product and took two pictures from the sides – one from the left at a 45-degree angle and one from the right at a -45-degree angle. Thirty total pictures were taken and tested. We observed that the accuracy of IQEngines service was 100 percent in the straight-on and left cases. However, in the right case, the accuracy was 80 percent, failing to recognize the product in two of the photos. The product recognition failures of these two photos occurred because of the following reasons. First, we took a picture of a product named "Hamburger Helper", specifically of the flavor "Chili Cheese". The IQ Engine service correctly recognized the bigger size of "Hamburger Helper", but incorrectly identified "Chili Cheese" as "Betty Crocker" instead. Second, we took a picture of a bottle of soy sauce from the right side. The IQEngine service did not read the information since the shape of the bottle, was cylindrical, causing the majority of text to wrap around the bottle, out of view. This bottle passed from the left, because more identifying features occur as part of the beginning of the product name, which is visible from the left. For these reasons, IQEngines shows worse performance from right-oriented photographs of grocery-store products, but a high accuracy of above 90 percent is still achievable overall.

5.4. Real-Grocery Subject Performance

We evaluated our application's in-person real-grocery-store functionality by analyzing the data we collected from the 15 in-person subjects: 87% from men and 13% from women. Participants' age ranged from 18 to 50, with the majority (53%) between 25 and 35 years of age. All results described here provide a comparison between the current visual guiding system being used at Lucky's Market (referred to as

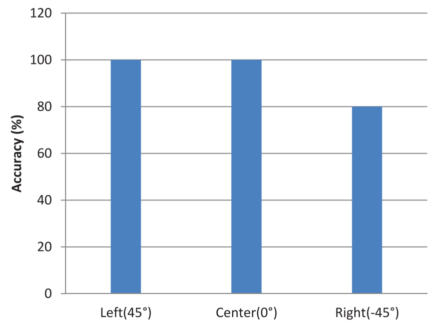

Fig. 17. IQEngines computer vision product identification and label return accuracy.

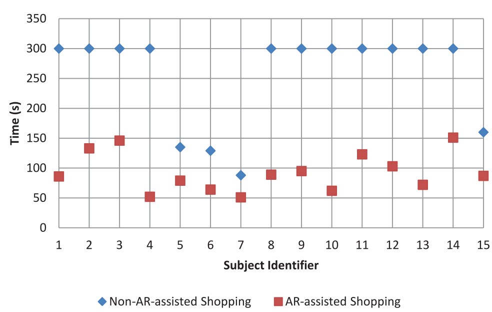

Fig. 18. Subject shopping speed without the constraint of ensuring the health of the selected product.

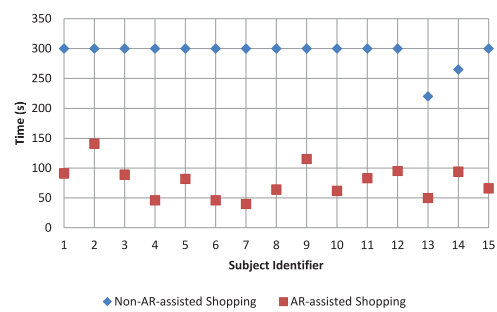

Fig. 19. Subject shopping speed under the constraint of ensuring the health of the selected product.

non-AR-assisted) and our AR-assisted smartphone app. Lucky's Market includes a two-level visual guiding system to direct users to the correct aisle and a customer service kiosk to receive nutritional information. In all our experiments, we did not influence the users in any way with respect to how they use the visual guiding system or the nutritional information kiosk.

In Figure 18, we conducted an experiment to measure the efficacy of the AR tags compared to non-AR-assisted. We asked users to find three products in the aisle, without regard to any health conditions. The time needed to find the three products with and without AR was compared. All AR tags were colored green. The experiment was set up so each individual was asked to find one set of three products without using our AR-assisted app and another set of three products using the AR tagging. Latency comparisons are therefore made across users rather than within the same user, since it would not be fair to ask a user to find the same three products by another method that they had just found. Figure 18 shows that for all 15 users that we tested, our AR-assisted tagging resulted in typically much faster performance 2X-3X in finding grocery store products. Most non-AR-assisted users in fact exceeded the maximum cap of 5 minutes that we set for the product discovery experiment, and would have taken longer in practice, so our 2X-3X estimate is a conservative one. Some users were quite savvy in using the list, but even in those cases the AR tagging results in faster discovery of recommended products.

In Figure 19, we conducted an experiment to measure the impact of healthy recommendations with and without AR tags. Again, we gave users a list of three products to find in the aisle, but in this case, one of the products was unhealthy. In the case of non-AR-assisted, a user may have to inspect the packaging, the nutrition facts label, or read through the ingredients in order to determine whether a product was unhealthy,

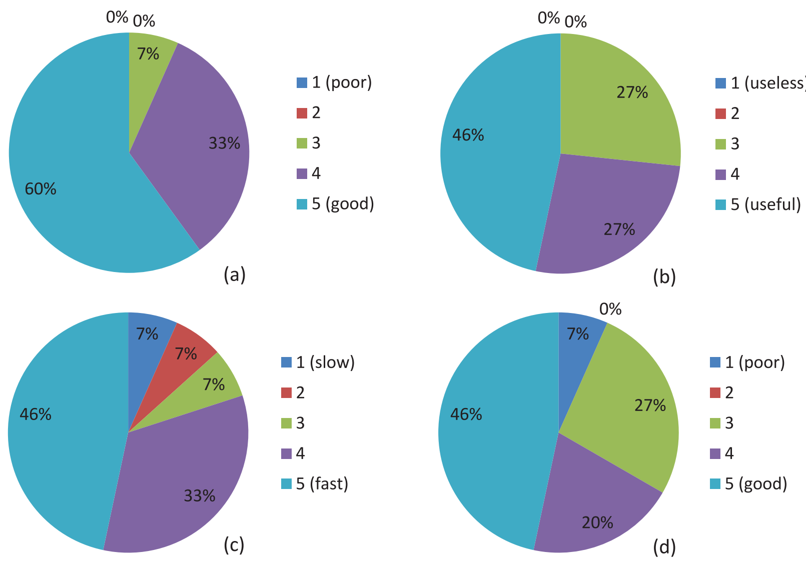

Fig. 20. In-person participants' satisfaction with aspects of our system: (a) Overall performance, (b) Usability, (c) Speed, (d) UI.

thus slowing down their shopping time. In contrast, our AR-based application already performs this filtering using our health recommendation subsystem on behalf of the user. Figure 19 shows that even the fast users from the earlier test are now so slowed by checking for healthy conditions that they are unable to finish within the 5-minute time limit, whereas in all cases the AR-assisted shopping finish in 2-1/2 minutes or less.

We also observe that our system remains similarly fast across both health-constrained and non-health-constrained shopping. Since the health-constrained test was performed after the health-free test, we hypothesize that users became more familiar with using our system, so the additional burden of ensuring that products are healthy was compensated for by increased familiarity with our mobile AR system.

Our test also examined the improvement our system provides in the identification of healthy grocery items, and conversely, the labeling and warning the user against purchasing products potentially unhealthy with respect to the specific user's dietary needs. We found that, when subjects did not use our system, they were actually able to correctly distinguish such healthy products from unhealthy products with perfect accuracy. Similarly, our system also performed with 100-percent healthy versus unhealthy identification. The improvement, however, came with the speed our system was able to do this versus the increase in time required for the subject to actually read the ingredient list themselves. Our system required no additional time.

Figure 20 shows the average satisfaction rating results. Almost all of the real grocery-shopping experiment participants (93%) were highly satisfied with our system's overall performance (5:60%, 4:33%) and the remaining 7% gave it a neutral satisfaction rating. About three-fourths (5:46%, 4:27%) of them were also highly satisfied with our application's usability, and the remaining one-fourth (3:27%) were neutrally satisfied. Participants' satisfaction with the speed of use of our system in enabling them to find healthy food products quickly was also rated quite highly by 79% of the participants (5:46%, 4:33%). The application's UI also received high satisfaction ratings from a large majority of the in-person experiment subjects, with two-thirds of them (5:46%, 4:20%) indicating they were highly satisfied with its UI.

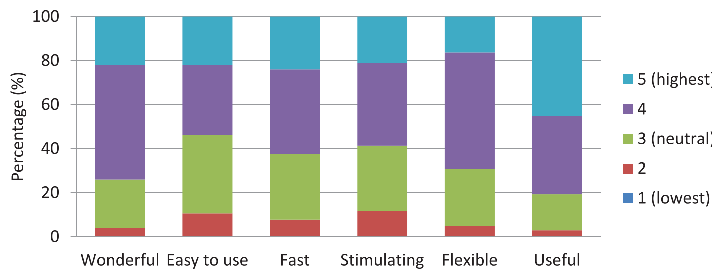

Fig. 21. Online participants overall satisfaction ratings of our system.

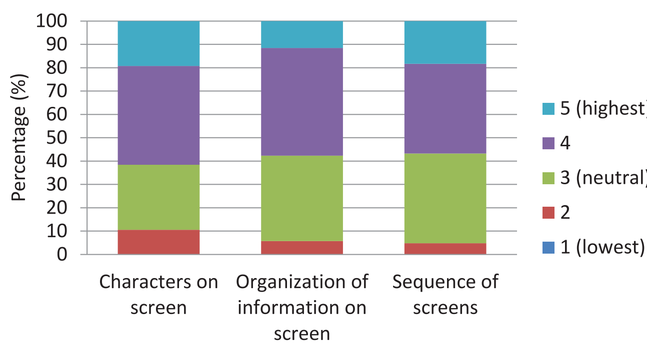

Fig. 22. Satisfaction ratings for the application screen interface.

5.5. On-Line Survey Evaluation of Our Application

Finally, we evaluated multiple features of our application based upon the QUIS (Questionnaire for User Interaction Satisfaction) [QUIS 1987] tool's structure, which is designed to assess users' subjective satisfaction with specific aspects of human-computer interfaces. As part of the survey reported in Section 3.1, the users were asked about their overall satisfaction and satisfaction with screen interface, and usability/UI.

Figure 21 shows the average rating results for different dimensions (overall-terrible:wonderful, difficult to use:easy to use, slow:fast, dull:stimulating, rigid:flexible, useless:useful) received from the 104 online participants. Overall, the participants were very satisfied with the features of our system. Three-quarters (74%: wonderful) of them were highly satisfied with the system overall and 80% of them indicated our system was very useful for the purpose it is intended. Over half to more than two-thirds of them reported it was: easy to use (54%), fast: (63%), stimulating (59%), and flexible: (69%). Only 3-11% of the participants rated these features unfavorably – giving them the low ratings.

Next, we asked participants to evaluate the screen interface: specifically the readability of the characters on the screen. Figure 22 shows average rating results. Almost two-thirds (62%) of the participants indicated that the characters displayed on the system's screen were easy to read. Three-fifths of the participants (58% and 57%, respectively) reported that the organization of information on screen was presented very clearly and the sequences of the screens presented was also quite clear and understandable. Only 5-10% of the participants rated these three features unfavorably.

Finally, we collected feedback from the participants on their impressions of the usability and UI features of our system. We asked them to rate the following features: use of colors and sounds, system feedback, system messages, and AR-tags, based on a 5-point "poor to good" scale. Figure 23 shows the participants' average satisfaction

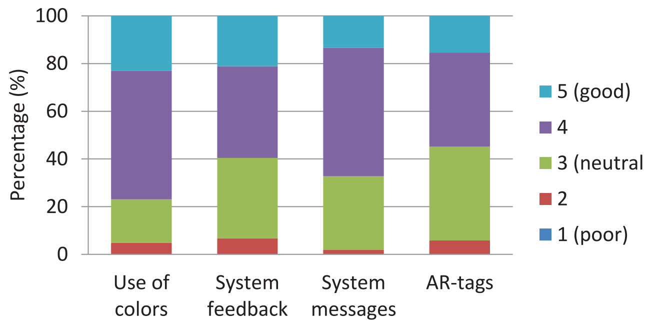

Fig. 23. Satisfaction ratings for usability and UI features.

rating results. The participants were quite highly satisfied with our system's usability and UI (Use of colors: 77%, System feedback: 60%, System messages: 67%, AR-tags: 55%). Only between 2 and 7% of the participants rated these features as functioning poorly (1/2).

In summary, a large majority of the 104 online survey participants who evaluated the video demonstration of our system were quite satisfied with its overall performance, screen interface, and usability/UI. In addition, feedback from the 15 in-person survey subjects who evaluated our system, after using it in a real grocery-store shopping experience, indicated that they were highly satisfied with its functionality. Taking both of these findings together into consideration, we expect that our system will prove to be very helpful to food shoppers who need to locate healthy food products in a grocery store quickly and efficiently.

5.6. Shopping-Based Personalized Pedometry

We evaluated the shopping-based pedometry algorithm, and focused on the accuracy of the number of footsteps measured with this algorithm, using the accelerometer sensor. When users are looking for products to purchase, they use our application by holding the phone vertically, pointed directly in front of them. Then when they identify a product to put in their cart or basket, they usually change the mobile phone's orientation as they approach the product – by either placing it on the cart handle, grasping it with the hand that is holding the basket handle, or by simply moving the hand that's holding the phone. We enhanced the pedometry algorithm to enable it to detect these shopping-based behaviors and to reduce the change-in-orientation errors. We performed an experiment to measure the accuracy of our algorithm's footstep detection, using the pedometry algorithm with five subjects, who were instructed to hold the phone in each of the 3 positions – on the cart handle, with the basket handle, and in their hand alone in a nonvertical position – as shown in Figure 24. Table II shows the average footsteps measured and the error rates obtained for the five subjects who used our application, while walking 100 steps for each of the three phone positions described. The shopping-based personalized pedometry closely determined the actual number of footsteps walked in each of the scenarios, with (a) 6.9%, (b) 5.7%, and (c) 4.5% error rates noted in the table. If we approximate each stride as about 3-feet long, then we're accumulating errors at a rate of 15 feet every 300 feet walked. The typical grocery store aisles that we tested in were about 40 feet long, and users did not spend 100 strides in a given aisle, so the accumulated error was small enough that it did not affect the perceived accuracy of overlaying of the AR tags within an aisle. However, missed strides will accumulate when considering whole-store cross-aisle navigation.

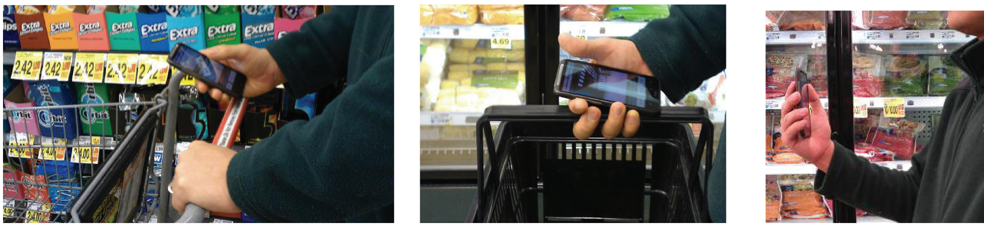

(a) (b) (c)

Fig. 24. Error rate of shopping-based personalized pedometry for three scenarios: (a) placing the phone on the cart handle, (b) holding the phone in hand with the basket handle, or (c) holding the phone in hand.

Table II. Footstep Detection Accuracy

User Status	Measured Avg. Steps	Actual Steps	Error (%)
With Cart	93.1	100	−6.9
With Basket	94.3	100	−5.7
Only with Hand	95.5	100	−4.5

6. CONCLUSION

This article has presented a mobile-based augmented reality system to help improve the ability of shoppers to find healthy food products in a grocery store aisle. We have shown that our application's color-based AR tagging functionality substantially reduces the amount of time it takes for shoppers to find desired healthy food products and avoid unhealthy ones. We conducted in-store evaluations of our system with 15 users of our application, and found that mobile AR tagging improved by at least 2X-3X the speed with which shoppers could find healthy products. We also conducted online surveys with over 100 subjects and found that 74% were highly satisfied with our application and only 3-11% were dissatisfied with our application.

REFERENCES

G. Agapito, B. Calabrese, I. Care, D. Falcone, P. H. Guzzi, N. Ielpo, T. Lamprinoudi, M. Milano, M. Simeoni, and M. Cannataro. 2014. Profiling basic health information of tourists: Towards a recommendation system for the adaptive delivery of medical certified nutrition contents. In *Proceedings of the 2014 International Conference on High Performance Computing Simulation (HPCS)*. 616–620. DOI:http://dx.doi.org/10.1109/HPCSim.2014.6903744

Junho Ahn, Mike Gartrell, James Williamson, and Richard Han. 2012. ARFusion: An indoor mobile augmented reality system supported by pedometry and context-awareness, an AR-assisted grocery shopping application. Institutional Review Board on the University of Colorado at Boulder 12-0102.

Junho Ahn and Richard Han. 2011. RescueMe: An indoor mobile augmented-reality evacuation system by personalized pedometry. In *Proceedings of the IEEE Asia-Pacific Services Computing Conference (APSCC)*.

Eileen Anderson, Richard Wientt, Janet Wojcik, Sheila Wientt, and Todd Bowden. 2001. A computerized social cognitive intervention for nutrition behavior: Direct and mediated effects on fat, fiber, fruits, and vegetables, self-efficacy, and outcome expectations among food shoppers. *Ann. Behav. Med.* 23, 2, 88–100.

Amy Barton, Lynn Gilbert, Julaluk Baramee, and Theresa Granger. 2006. Cardiovascular risk in hispanic and non-hispanic preschoolers. *National Institute of Health* (May-June 2006).

Stephane Beauregard. 2007. Omnidirectional pedestrian navigation for first responders. In *Proceedings of the 4th Workshop on Positioning, Navigation and Communication (WPNC'07)*.

J. Bobadilla, F. Ortega, A. Hernando, and A. GutiéRrez. 2013. Recommender systems survey. *Know.-Based Syst.* 46 (2013), 109–132. DOI:http://dx.doi.org/10.1016/j.knosys.2013.03.012

Cristina Botella, Juani Breton-Lopez, Soledad Quero, Rosa Banos, Azucena Garcia-Palacios, Irene Zaragoza, and Alcaniz Raya. 2011. Treating cockroach phobia using a serious game on a mobile phone and augmented reality exposure: A single case study. *Comput. Hum. Behav.* 27, 1, 217–227.

DanKam. 2010. DanKam AR application for color blindness. http://news.cnet.com/8301-27080_3-20026054-245.html.

DietGuideline 2010. U.S. Department of Agriculture. 2010 Dietary guidelines for Americans. http://www.cnpp.usda.gov/DGAs2010-PolicyDocument.htm. (2010).

Jill Freyne and Shlomo Berkovsky. 2010. Intelligent food planning: Personalized recipe recommendation. In *Proceedings of the 15th International Conference on Intelligent User Interfaces (IUI'10)*. 321–324. DOI:http://dx.doi.org/10.1145/1719970.1720021

Christoph Fuchs, Nils Aschenbruck, Peter Martini, and Monika Wieneke. 2011. A survey on indoor tracking for mission critical scenarios. *Perva. Mobile Comput.* 7, 1.

Subhashini Ganapathy, Glen Anderson, and Igor Kozintsev. 2011. MAR shopping assistant usage: Delay, error, and utility. In *Proceedings of the IEEE Virtual Reality Conference (VR)*. 207–208.

Golfscape. 2014. Golfscape GPS AR range finder. http://golfscapeapp.com/.

GoogleDocs 2014. Google Docs' survey tool. http://support.google.com/docs/bin/answer.py?hl=en&answer=87809.

Levent Görgü, Abraham Campbell, Kealan McCusker, Mauro Dragone, Michael O'Grady, Noel O'Connor, and Greg O'Hare. 2010. FreeGaming: Mobile, collaborative, adaptive and augmented exergaming. In *Proceedings of the 8th International Conference on Advances in Mobile Computing and Multimedia (MoMM'10)*. 173–179.

Yanying Gu, Anthony Lo, and Ignas Niemegeers. 2009. A survey of indoor positioning systems for wireless personal networks. *IEEE Commun. Surv. Tutori.* 11, 1, 13–32.

Ramón Hervás, Alberto Garcia-Lillo, and José Bravo. 2011. Mobile augmented reality based on the semantic web applied to ambient assisted living. In *Ambient Assisted Living*, Jos Bravo, Ramn Hervs, and Vladimir Villarreal (Eds.), Lecture Notes in Computer Science, vol. 6693, Springer Berlin / Heidelberg, 17–24.

Hwajung Hong, Hee Young Jeong, Rosa I. Arriaga, and Gregory D. Abowd. 2010. TriggerHunter: Designing an educational game for families with asthmatic children. In *Proceedings of CHI EA'10*. 3577–3582.

IQEngines. 2013. IQEngines: Image recognition and visual search. http://www.iqengines.com.

Vaiva Kalnikaite, Yvonne Rogers, Jon Bird, Nicolas Villar, Khaled Bachour, Stephen Payne, Peter Todd, Johannes Schoning, Antonio Kruger, and Stefan Kreitmayer. 2011. How to nudge in Situ: Designing Lambent devices to deliver salience information in supermarkets. In *Proceedings of UbiComp'11*.

David Katz, Valentine Njike, Zubaida Faridi, Lauren Rhee, Rebecca Reeves, and David Jenkins. 2009. The stratification of foods on the basis of overall nutritional quality: The overall nutritional quality index. *Amer. J. Health Promot.* (Nov.–Dec. 2009).

Stefan Ladstaetter, Patrick Luley, Alexander Almer, and Lucas Paletta. 2010. Multisensor data fusion for high accuracy positioning on mobile phones. In *Proceedings of MobileHCI*. 395–396.

Jia-Kuan Lin, Po-Hsun Cheng, Yen Su, et al. 2011. Augmented reality serious game framework for rehabilitation with personal health records. In *Proceedings of the 13th IEEE International Conference on e-Health Networking Applications and Services (Healthcom)*. IEEE, 197–200.

T. Lobstein and S. Davies. 2009. Defining and labeling 'healthy' and 'unhealthy' food. *Public Health Nutr.* (2009).

Jennifer Mankoff, Gary Hsieh, Ho Chak Hung, Sharon Lee, and Elizabeth Nitao. 2002. Using low-cost sensing to support nutritional awareness. In *Proceedings of the 4th International Conference on Ubiquitous Computing*. Springer-Verlag, 371–376.

MechanicalTurk 2014. Amazon Mechanical Turk Survey Services. https://www.mturk.com/mturk/.

Cliona Ni Mhurchu, Tony Blakely, Joanne Wall, Anthony Rodgers, Yannan Jiang, and Jenny Wilton. 2007. Strategies to promote healthier food purchases: A pilot supermarket intervention study. *Public Health Nut.* 10, 06, 608–615.

Barry Mulrooney, Mairéad McDermott, and Nick Earley. 2006. NutraStick: Portable diet assistant. In *CHI'06 extended abstracts on Human factors in computing systems (CHI EA'06)*. ACM, New York, 1855–1860. DOI:http://dx.doi.org/10.1145/1125451.1125802

Yoosoo Oh, Ahyoung Choi, and Woontack Woo. 2010. u-BabSang: A context-aware food recommendation system. *J. Supercomput.* 54, 1, 61–81. DOI:http://dx.doi.org/10.1007/s11227-009-0314-5

PCAST. 2010. Realizing the full potential of health information technology to improve healthcare for americans: The path forward. http://www.whitehouse.gov/sites/default/files/microsites/ostp/pcast-health-it-report.pdf. (2010).

M. Phanich, P. Pholkul, and S. Phimoltares. 2010. Food recommendation system using clustering analysis for diabetic patients. In *Proceedings of the International Conference on Information Science and Applications (ICISA)*. 1–8. DOI:http://dx.doi.org/10.1109/ICISA.2010.5480416

QUIS 1987. Questionnaire for user interaction satisfaction. http://lap.umd.edu/quis/about.html.

SkyMap. 2011. Google Sky Map AR astronomy application. http://www.google.com/mobile/skymap/.

Richard Wientt, Eileen Smith Anderson-Bill, Patricia G. Bickley, Janet Walberg-Rankin, John F. Moore, and Michael Leahy. 1997. Nutrition for a Lifetime System ©: A multimedia system for altering food supermarket shoppers' purchases to meet nutritional guidelines. *Comput. Hum. Behav.* 13, 3, 371–392. DOI:http://dx.doi.org/DOI: 10.1016/S0747-5632(97)00015-0

WordLens. 2014. WordLens augmented reality language translation app for the iPhone. http://www.questvisual.com/.

Received January 2015; revised April 2015; accepted June 2015