LiveCompare: Grocery Bargain Hunting Through Participatory Sensing

Linda Deng
Duke University
Durham, NC
linda@cs.duke.edu

Landon P. Cox
Duke University
Durham, NC
lpcox@cs.duke.edu

ABSTRACT
Many consumers are misled into paying high prices due to the search costs associated with attaining price information [16]. The popularity of bargain-hunting web sites like Slickdeals.net, which boasts 2.5 million visitors per month, hints that many shoppers are indeed in search of tools to help them save money. We present LiveCompare, a system that leverages the ubiquity of mobile camera phones to allow for grocery bargain hunting through participatory sensing. We utilize two-dimensional barcode decoding to automatically identify grocery products, as well as localization techniques to automatically pinpoint store locations. We show that an incentive scheme is inherently ingrained into our query/response protocol, and we suggest self-regulating mechanisms for preserving data integrity. As a result, we demonstrate that money-saving price comparisons can be conducted among brick and mortar grocery stores without the explicit cooperation of the stores.

1. INTRODUCTION
Despite the pervasiveness of online shopping sites, many people still prefer to purchase everyday grocery and household items from brick and mortar stores [11]. However, identical products often sell for varying prices among different stores [16]. For example, on a given day in Durham, North Carolina, we observed that the same brand of toilet paper cost twice as much at a CVS drugstore as at a nearby Harris Teeter supermarket.

Such price dispersion exists partly because of the high search cost of obtaining price information [16]. For many products like books and electronics, existing online services like Google Product Search [9] can provide useful pricing information for popular online merchants (e.g., Amazon.com) and even give a good indication of brick and mortar prices (e.g., by presenting the price listed at Walmart.com). However, it remains difficult to find online pricing information for grocery items not commonly purchased online, such as a carton of ice cream. Although some grocery stores offer online shopping options and post their prices online, most merchants do not. Additionally, although our local Harris Teeter store provides an online ordering service that divulges all current in-store prices, a user agreement on the web site prohibits use of the data for any purpose other than online shopping; thus, it would not be possible to utilize this data within a price comparison application.

To facilitate price comparison of grocery items, many deal-seeking consumers participate in online forums like AFullCup.com to query other users for "price checks" on their items of interest at particular stores. Unfortunately, the pricing information gathered via this method is difficult to parse and organize, subject to human errors, and not readily accessible to users already located inside a grocery store.

With the increasing ubiquity of mobile devices, mobile applications like CompareEverywhere [12] and ShopSavvy [2] aim to cater toward on-the-fly price comparisons. However, these applications rely on online information as well, thus providing an incomplete picture, as they are unable to determine price information for grocery items that are not sold online. Another mobile application, MobiShop [4], facilitates price comparisons among grocery items but utilizes error-prone optical character recognition (OCR) techniques and relies on receipt information, which is difficult for a computer to parse.

In this paper, we propose LiveCompare, a system based on *participatory sensing* with mobile devices to improve inter-store grocery price comparisons. LiveCompare participants use their camera phones to snap a photograph of the price tag of their product of interest. The product is uniquely identified via a barcode included on the price tag in most grocery stores. The photograph is then uploaded to a central repository for satisfying future queries. In exchange for submitting this price data point, the user receives pricing information for the scanned product at other nearby grocery stores.

As with other participatory sensing applications, LiveCompare faces two inherent challenges. The first challenge is to give users an incentive to contribute data. LiveCompare's exchange-based structure is uniquely well-suited to meeting this challenge: users only benefit from the application when the server *compares* their submitted price to other related data points. As a result, the data pool can only be queried as users simultaneously contribute.

HotMobile 2009, February 23–24, 2009, Santa Cruz, CA, USA.

The second challenge is to ensure data integrity. LiveCompare provides high-quality data through two complementary social mechanisms. First, LiveCompare relies on humans, rather than machines, to interpret complex sale and pricing information. The only part of a price tag that must be interpreted by a computer is the barcode to retrieve a Universal Product Code (UPC). Because of this, LiveCompare does not need to rely on error-prone OCR algorithms to extract textual tokens or on linguistic models to make sense of sets of tokens. Second, each LiveCompare query returns a subset of the data pool for a user to consider. If an image does not seem relevant, the user can quickly flag it. This allows users to collectively identify malicious data.

The rest of this paper is organized as follows. We discuss related work in more detail in Section 2. Section 3 presents the architecture of LiveCompare, along with proposed solutions to the challenges of incentives and data integrity. In Section 4, we discuss the findings of our field work in various local grocery stores. Finally, we conclude in Section 5.

2. RELATED WORK

The ubiquity of camera phones has enabled many applications that use photographs to retrieve information. Color Match [10] is an application in which consumers transmit a photograph of their face and a reference color chart to a server that normalizes the colors within the photograph and determines the user's optimal foundation makeup shade; thus, users can receive immediate feedback on which product to purchase when standing in the beauty aisle at a store. Point&Find [6] is an augmented reality application through which users can retrieve information about objects they photograph, such as points of interest in tourist settings; GPS data is used to restrict the search space for conducting image recognition of the point of interest. In both Color Match and Point&Find, the data used to satisfy queries is collected independently by the service provider.

Participatory sensing [5] envisions the use of ubiquitous mobile devices to enable the collection and sharing of "local knowledge" for applications in areas such as public health, urban planning, cultural analysis, and natural resource management. Micro-Blog [8] is a general-purpose platform for the sharing of geotagged multimedia "blogs." Unlike LiveCompare, MicroBlog does not offer users with incentives to share.

The use of mobile devices to facilitate bargain hunting while browsing brick and mortar stores is not a new idea. In the late nineties, the Pocket BargainFinder [3] sought to provide consumers with the ability to scan barcodes in retail stores to perform on-the-fly price comparisons. Two more recent examples include $275,000 winners of Google's Android Developer Challenge [1]—CompareEverywhere [12] and ShopSavvy [2]—which aim to operate on consumer camera phones (rather than application-specific hardware like the Pocket BargainFinder). However, similar to the Pocket BargainFinder, CompareEverywhere and ShopSavvy rely on crawling existing online databases to determine pricing information and are thus not suitable for grocery items not commonly purchased online.

Finally, a recent application, MobiShop [4], is a price comparison system leveraging participatory sensing through geotagged receipt scanning. It differs from LiveCompare in two important ways. First, LiveCompare utilizes individual price tag scanning to avoid the error-prone parsing of receipt information. For example, often a product will be listed on a receipt with a description like "CHICKEN STRIPS" without any indication of size or brand; such items would be difficult to insert into a database in an easy-to-query format. Secondly, LiveCompare offers explicit incentive and integrity mechanisms to encourage participation.

3. DESIGN

LiveCompare assumes that each participating grocery-store shopper has a camera phone with Internet access. When a user finds a product of interest, he snaps a photograph of the product's price tag. From the photograph, the user's phone extracts information about the product using the unique UPC barcode located on the tag. In most grocery stores, price-tag barcodes are identical to the barcodes on the actual products and facilitate global product identification. Decoding a barcode from a photograph can be performed quickly using barcode libraries such as ZXing [17]. One caveat is that photographs must be reasonably clear at close focus—achievable with any camera phone that supports a macro mode, such as the Nokia N95.

LiveCompare relies on the mobile device to interpret barcodes. This allows the device to extract barcodes from a local high-quality image and to transfer a smaller lower-quality image to the server for better network performance. Once the barcode has been decoded on the device, the numerical UPC value and the just-taken photograph are transferred to LiveCompare's central server. This data is stored in LiveCompare's database for use in future queries, and the UPC value determines the unique product for which price comparisons are requested. The client also sends its GPS or GSM cell information to the server so that the current store can be identified. This location information allows LiveCompare to limit query results to include only nearby stores (where a distance threshold may be set by the user). Results include store information and the option to view the timestamped photographs associated with the specific product in question at each store.

An important feature of LiveCompare is that it only presents users with raw photographic data collected by other users, rather than attempt to automatically extract numerical pricing information. This is because techniques such as optical character recognition (OCR) are error-prone, slow, and also may have difficulty extracting the true price from a tag containing multiple sale-price values or a special offer such as "buy one, get one free." Additionally, users are not required to manually input pricing data; this improves data integrity and user convenience.

LiveCompare's low burden of participation provides a clear path to deployment. All items are identified via a globally unique product code that is already widely used. Additionally, a central authority need not worry about the collection and maintenance of data, since data is gathered every time a user initiates a query. The primary drawback of LiveCompare's architecture is that it does not scale down well—i.e., when few queries have been submitted, query results are unlikely to be helpful. To ameliorate this issue, LiveCompare can fall back on existing price pools such as online grocery stores and drugstores to provide comparisons when data is scarce. Pricing information from online stores will be lower quality than data contributed by LiveCompare users, but it still provides a valuable reference point to help consumers determine whether a particular item is reasonably priced.

3.1 Incentives

In many participatory sensing systems, incentives are a fundamental challenge. To contribute data, users give up their time, attention, and mobile device's battery power.

LiveCompare directly addresses this challenge through its query protocol. When a user submits a query from a grocery store, he identifies the product for which he wants price comparison information by snapping a photograph of the product's price tag (which includes a unique UPC barcode). This also allows the server to append additional information to its database (i.e., the pricing information for the product, which is included in the price tag, as well as the physical location of the user determined via localization techniques). Thus, by requiring that a geotagged photograph be uploaded as part of a query, LiveCompare automatically populates its database whenever a user initiates a query.

3.2 Integrity

LiveCompare reduces the risk of operator error by requiring product identification and pricing information to be submitted in the same photograph. Such tight coupling reduces the likelihood that a user will inadvertently input an incorrect price. Retrieving product information through barcode decoding and identifying store information through device localization further decrease the risk of human error.

In the case of a malicious user who crafts his own price tags to submit to the database, LiveCompare can employ a cooperative anomaly-detection effort. If LiveCompare users encounter an entry that is clearly inauthentic, they can flag the entry so that it can be moderated or automatically removed when flagged by a sufficient number of other users. If a particular user is found to be the owner of several suspicious entries, he may be banned. A similar tagging system can also be employed to correct for computer errors such as incorrect barcode decoding (although the offending participant may not be penalized so severely for such errors).

Unfortunately, there is no explicit incentive for users to collaborate in such a community effort to maintain data integrity. Nonetheless, community-driven web sites such as Wikipedia [14] have shown that people will often voluntarily contribute for no monetary compensation. A reputation system could also be utilized to encourage users to submit useful data. Point systems like the one employed by Yahoo! Answers [15] indicate that many users value reputation points, even if they possess no material meaning.

Another possible approach is to perform OCR to ensure that different photographs submitted by different users—but for the same product at the same store—contain similar extracted tokens. If a photograph contains wildly different tokens from other photographs for the same product/store, LiveCompare can automatically flag the entry as suspect.

3.3 Limitations

Since LiveCompare relies on UPCs as globally unique keys, generic-brand products cannot be easily compared across different stores. However, a UPC is structured such that five known digits denote the manufacturer code; thus, these five digits would be the same for every item within a store's line of generic-brand products. It would therefore be easy to determine via the UPC that a product is in fact generic-brand. Unfortunately, it would be non-trivial to then determine the identity of the actual product so that it may be compared against the same product of another generic brand. To accomplish this, LiveCompare may need to resort to performing OCR on the price tag and/or requiring human input to identify the product.

Occasionally a price tag may fail to include a barcode; for example, we have observed this on some clearance items with handwritten price tags, as well as several items in the produce and deli departments. In such cases, if the product itself contains a UPC, the user may take a photograph of the product's UPC next to the corresponding price tag. On the other hand, if the product does not have a UPC at all (e.g., produce sold by weight), OCR or human input may again be required. In either case, data integrity may be compromised, but it is our hope that self-regulation would keep malicious manipulation at bay.

Finally, we have not made privacy a primary concern in the design of LiveCompare. However, users can create a LiveCompare-specific pseudonym to interact with the application that need not have any relationship to their real identity. To prevent the identification of users via IP address, an anonymity network like Tor [7] can be used.

4. EVALUATION

To determine the usefulness and feasibility of LiveCompare, we conducted field work in seven different brick and mortar stores throughout Durham, North Carolina. We show that price dispersion can be observed across a variety of grocery items, that typical store price tags contain sufficient information to enable LiveCompare's infrastructure, and that data transfer performance is reasonable over an HSDPA network. We also propose some localization methods that provide a fine enough granularity to detect individual stores.

4.1 Price dispersion

To demonstrate that price dispersion exists across a range of products commonly purchased at brick and mortar locations, we traveled to seven different stores around Durham. Using a list of ten common grocery items, we photographed price tags of some of these items at each store. Specifically, we visited three conventional grocery stores (Food Lion, Harris Teeter, and Kroger), an organic foods store (Whole Foods Market), a hypermarket (SuperTarget), a discount store (Kmart), and a drugstore (CVS); all seven stores were located within a six-mile diameter. These visits all occurred on the same day (October 5, 2008), and we compared only items that were identical (i.e., the same size and variety) and had identical UPCs across stores. Table 1 depicts our findings, which show that price dispersion exists across a variety of products and stores. In cases in which an item was on sale, we considered only the current sale price on October

Table 1: Price ranges of 10 grocery items, each found at 3-5 local stores on October 5, 2008.

Item	Price range	Low price store	High price store	Other stores
Ben & Jerry's ice cream	$3.00-$4.49	Food Lion	Harris Teeter	Target, Whole Foods
Coca-Cola soft drink	$1.11-$1.59	Harris Teeter	Kroger	Kmart, Target
Colgate toothpaste	$3.99-$4.99	Target	Harris Teeter	CVS
Cottonelle toilet paper	$5.99-$11.99	Harris Teeter	CVS	Target
Gillette power razor	$7.94-$11.99	Target	CVS	Harris Teeter
Herbal Essences shampoo	$2.49-$3.79	Food Lion	Harris Teeter, Kmart	CVS, Target
Kashi cereal	$2.66-$4.59	Target	CVS	Food Lion, Kroger, Whole Foods
Kraft cheese slices	$3.59-$4.69	Target	Kroger	Harris Teeter
Tide laundry detergent	$10.00-$16.49	Target	CVS	Food Lion, Kmart
Tropicana orange juice	$2.99-$3.99	Kroger, Target	Whole Foods	Food Lion, Harris Teeter

(a) Harris Teeter (b) Kroger (c) Target

Figure 1: Example photographs of price tags for Kraft cheese slices, taken with N95 in (a) Harris Teeter, (b) Kroger, and (c) Target.

5, 2008, rather than the regular price—as such sales often contribute toward significant price dispersion.

Although certain patterns can be discerned from this data (e.g., that Target often has the lowest price, or that CVS often has the highest price), no single store *always* has the lowest or highest price for all items; indeed, Harris Teeter simultaneously offers the lowest price on some items and the highest price on others! Additionally, price differences are often significant: for example, a twelve-pack of double rolls of Cottonelle toilet paper was twice as expensive at CVS as at Harris Teeter. Incidentally, the particular CVS we visited is even located in the same shopping complex as Harris Teeter, and the two stores are easily within walking distance of each other. In a case like this, it is clear that a toilet-paper-seeking CVS shopper could potentially derive great savings through the use of LiveCompare.

4.2 Price tag content

Figure 1 shows the photographs of three price tags that we captured with a Nokia N95 8GB camera phone. Note that each price tag includes a barcode corresponding to the globally identifying UPC of the product, while the human-readable descriptions are not easy to parse (e.g., it may be unclear that "KRAFT STK PK AMER SNGLS-C" indeed refers to cheese slices). This reinforced our intuition that barcodes are a more convenient method of global and unique identification than OCR.

We additionally verified that the barcodes in our price tag photographs could be decoded via a barcode library such

Figure 2: Example photograph of a CVS price tag, which includes detailed deal information that would be difficult to glean from a receipt.

as ZXing [17], even when the barcode was only a fraction of the noisy photograph. In doing this, we discovered an anomaly: many of the barcodes that appear on the price tags in CVS do not actually correspond to a UPC. We assume that CVS utilizes an internal identification scheme; although the barcodes can be decoded, it may be difficult to map them onto the same product at a different store chain. This is also a problem with generic-brand products whose unique identifiers exist only within one particular store chain. We intend to further explore ways to identify such products,

Table 2: HSDPA transfer rates for uploading 18.3 KB and downloading 71.3 KB across 20 trials.

	Average	Standard deviation
Uploading 18.3 KB image	4.08 seconds	1.04 seconds
Downloading 71.3 KB in images	3.57 seconds	1.39 seconds
Total latency of upload/download operation	7.65 seconds	1.88 seconds

as well as certain produce and deli items that often do not contain any barcodes on the price tag. In such cases, OCR or human input may be required.

Of course, each price tag also contains the product's actual pricing information, sometimes including effective sale dates and whether a shopper's card may be required to take advantage of the sale price. This information is relevant to a querying user, and thus providing a snapshot of the price tag is generally sufficient as a response to a LiveCompare price comparison query. Additionally, the information provided on a price tag is often more useful than what can be gleaned from a receipt (as used in MobiShop [4]).

As an example, Figure 2 depicts a CVS price tag promoting a sale price and purchase reward that are both dependent on the customer utilizing a CVS shopper's card. From this photograph, a user would be able to discern the CVS card price of $3.97, the $2 reward that would be provided within CVS's loyalty program upon purchase of the product, and even the effective sale dates (we verified that the small font describing the sale dates is readable even on the Nokia N95's screen). In contrast, it is much more difficult to extract this information from the register receipt attained when purchasing the product. First, on the receipt, the price is conveyed via the following line:

```
1 GE CHRRY TNG  1L      3.97B SAVED      .02
```

Nine inches below this line on the receipt, the following line finally conveys the purchase reward associated with the product:

```
Sparkling Water,Buy 1 Get 2 EB
```

Although both the purchase price and reward amount are relevant in determining the desirability of the deal, the spatial separation of these two pieces of information, along with the difficulty of equating "GE CHRRY TNG 1L" with "Sparkling Water," render it nearly impossible to attain the same level of information from the receipt as from the price tag in Figure 2. In addition, since price tags can be scanned without a user actually purchasing the product, LiveCompare's method of price capture facilitates a larger database than a system relying on receipt scanning alone.

4.3 Data transfer

We measured the speed of data transfer over AT&T's HSDPA network to demonstrate the latency of a typical LiveCompare query. We note that if barcode decoding is performed on the mobile device, the photograph to be uploaded to LiveCompare's central server need not actually have a high resolution; it would merely need to be large enough to be comfortably viewable by other users on their mobile device screens. We also note that on many Symbian devices such as the Nokia N95 8GB, three differently sized thumbnails are automatically generated inside a hidden folder on the file system whenever a photograph is taken. Thus, instead of incurring extra compression overhead, we observe that the largest of these thumbnails (320 $\times$ 320 resolution) is sufficient for viewing price tag information on a mobile device screen; we therefore choose to send this image to LiveCompare's central server.

In our data transfer experiment, we transmitted one of these 320 $\times$ 320 thumbnails to our server while standing inside our local Whole Foods grocery store; this image was 18.3 KB in size. Our server then responded to the query with four image files of price tag photographs taken at other grocery stores; each image was sized between 14 KB and 22 KB, and the total size of transferred images was 71.3 KB. Table 2 depicts the data rates we achieved, showing that a typical query may be serviceable with only a few seconds of delay.

4.4 Localization

We also evaluated the feasibility of GPS, Wi-Fi, and GSM localization when visiting our seven local stores. Assisted GPS was able to quickly attain accurate coordinates for all of the stores in either the checkout area or just outside the entrance of the store; however, GPS was useless when deep within the aisles of the stores, since satellite signals could not be sufficiently detected indoors. Wi-Fi localization was also impossible in one of the stores we visited, since no access points could be detected. Additionally, it appears that many parts of Durham have not been sufficiently wardriven for us to rely on popular Wi-Fi localization databases such as that provided by Skyhook Wireless [13]. Finally, simply identifying the current GSM cell was also insufficient for localization at the granularity of individual stores; in our experiment, four of the seven stores we visited were actually located in two different shopping complexes across the street from each other, and the GSM cell did not change when moving among these four stores.

As a result, it is unlikely that any one of these schemes by itself could provide the fine-grained localization that LiveCompare needs to automatically identify the user's current store. Since power constraints prevent users from always maintaining their GPS coordinates, a small amount of user assistance may be required for localization. GPS localization could thus be attempted initially when a user first launches LiveCompare in or near a particular store; then, if sufficient satellite signal cannot be detected, the user could be prompted to resolve any ambiguities associated with identifying the store via GSM cell. (Future temporally proximate queries from the same GSM cell could then be assumed to be from the same store.) Alternatively, the use of hybrid always-on energy-efficient localization schemes, such as the one described in [8], could also be explored.

5. CONCLUSIONS

In this position paper, we presented LiveCompare, a system to enable price comparison of grocery items through participatory sensing. We suggested the combined use of barcode decoding and GPS/GSM localization to automate the detection of product identity and store location. We also discussed a novel incentive scheme to encourage participation and proposed methods for maintaining data integrity. Finally, we presented real-world examples of price dispersion to show that LiveCompare can be used as a convenient way for everyday grocery shoppers to save money.

6. REFERENCES

[1] Android Developer Challenge. http://code.google.com/android/adc.html.

[2] R. Barnes. ShopSavvy. http://www.biggu.com.

[3] A. B. Brody and E. J. Gottsman. Pocket BargainFinder: a handheld device for augmented commerce. In *Handheld and Ubiquitous Computing*, 1999.

[4] N. Bulusu, C. T. Chou, S. Kanhere, Y. Dong, S. Sehgal, D. Sullivan, and L. Blazeski. Participatory sensing in commerce: using mobile camera phones to track market price dispersion. In *UrbanSense*, 2008.

[5] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava. Participatory sensing. In *WSW at Sensys*, 2006.

[6] G. Cuellar, D. Eckles, and M. Spasojevic. Photos for information: a field study of cameraphone computer vision interactions in tourism. In *CHI Extended Abstracts on Human Factors in Computing Systems*, 2008.

[7] R. Dingledine, N. Mathewson, and P. Syverson. Tor: the second-generation onion router. In *USENIX Security Symposium*, 2004.

[8] S. Gaonkar, J. Li, R. R. Choudhury, L. Cox, and A. Schmidt. Micro-Blog: sharing and querying content through mobile phones and social participation. In *MobiSys*, 2008.

[9] Google Product Search. http://www.google.com/products.

[10] J. Jain, N. Bhatti, H. H. Baker, H. Chao, M. Dekhil, M. Harville, N. Lyons, J. Schettino, and S. Süsstrunk. Color Match: an imaging based mobile cosmetics advisory service. In *Mobile HCI*, 2008.

[11] R. K. Rajamma, A. K. Paswan, and G. Ganesh. Services purchased at brick and mortar versus online stores, and shopping motivation. In *Journal of Services Marketing*, 2005.

[12] J. Sharkey. CompareEverywhere. http://compare-everywhere.com.

[13] Skyhook Wireless. http://www.skyhookwireless.com.

[14] Wikipedia. http://www.wikipedia.org.

[15] Yahoo! Answers. http://answers.yahoo.com.

[16] Y. Zhao. Price dispersion in the grocery market. In *The Journal of Business*, 2006.

[17] ZXing. http://code.google.com/p/zxing.