Physical Analytics: A New Frontier for (Indoor) Location Research

Rajalakshmi Nandakumar
Swati Rallapalli†
Krishna Chintalapudi
Venkata N. Padmanabhan
Lili Qiu†
Aishwarya Ganesan
Saikat Guha
Deepanker Aggarwal
Aakash Goenka

Microsoft Research India

†*The University of Texas at Austin*

October 2013

Technical Report
MSR-TR-2013-107

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Physical Analytics: A New Frontier for (Indoor) Location Research

R. Nandakumar, S. Rallapalli[†], K. Chintalapudi, V. N. Padmanabhan, L. Qiu[†], A. Ganesan, S. Guha, D. Aggarwal, and A. Goenka[*]
Microsoft Research India [†]University of Texas at Austin

ABSTRACT

The large body of research on localization has been driven by the goal of providing users with location-based services, be it mapping and navigation or local alerts and advertisements. However, as we discuss in this paper, location information can over time provide deep insight about about a user, going well beyond the location domain itself. We argue that unlocking the wealth of information available from location is valuable and represents a new and promising frontier for location-related research, especially in the indoor domain. We call this Physical Analytics, analogous to Online Analytics, with footsteps taking the place of a clickstream. We describe research opportunities, challenges, and our initial investigation in Physical Analytics.

1. INTRODUCTION

Localization technologies and their applications have received much attention, both in the research community and increasingly in the commercial world. The focus has largely been on using location information to provide location-based services to users, e.g., mapping, navigation, alerts, reminders, advertisements, etc. While such services tied to location have been very useful in the outdoor setting (e.g., GPS navigation), the search of a similar, user-facing killer app for location-based services indoors is a continuing quest.

We argue that analyzing user location information in the physical world — *Physical Analytics, or Phytics* — to gain deep insight about users, is a promising new frontier for location-related research. Just as a clickstream in the online world provides insights about users, so does footsteps in the physical world. We focus primarily, though not exclusively, on physical analytics in indoor settings, motivated by the observation that people spend much of their time indoors and have much of their interactions with each other and with businesses indoors.

We can draw many parallels between the online and the physical worlds. Analogous to the online world, analytics in the physical setting could be of value to users, advertisers, and venue owners (who can be thought of as the analogs of "publishers" in the online world). Users would benefit from personalized information, advertisers from more effective targeting, and venue owners from insights that enables them to deploy their (physical) real estate in a manner that helps maximize engagement with users. In the online setting, how users traverse web pages and applications, including the links they click, the search keywords they type, where they dwell (i.e., one which websites and apps), and where their mice point to are indicative of the user's profile and interests. Analogously, the traversal made by a user in the physical world, including where they go, how long they dwell, and what interactions they have (e.g., picking up items from a store shelf) reveals a lot about them. Similar analogies can be drawn between the online and physical worlds when considering the behavior of a population of users instead of just individual users.

However, the physical world is also different from the online one in significant ways that present interesting challenges and opportunities in the context of Physical Analytics. In the online setting, an HTTP transaction resulting from a click encapsulates all the information needed to track the user, using mechanisms such as HTTP redirection, cookies, and tracking pixels. This makes tracking convenient and also makes it easy to interpose third-party services for the purposes of gleaning information about the user. In contrast, tracking in the physical world is much more challenging because the necessary information is incomplete and also scattered. For example, the user's mobile device (e.g., smartphone) might be in a position to track the user's every movement but might lack both accuracy and semantic context with regard to locations (e.g., knowing that the user's location corresponds to the women's clothing section in a retail store). The venue owner would have the context information but may not be in a position to track the user, especially outside the confines of their physical space. A third party such as a cellular network operator might have the footprint to have broad visibility but is unlikely to have fine-grained or semantic information.

On the other hand, since everything on the web is virtually a click away, a user's pattern of traversal online could be quite complex, so the relative ease of online tracking is indeed a boon for taming this complexity. In the physical world, however, traversal is constrained by physical location and proximity. Such constraints can be used to fill in the holes with regard to tracking. For example, if over time a user is never seen at the entrance of a mall that houses the only Apple Store in town, we can safely conclude that the user has not been visiting the Apple Store, even if we do not have any visibility into locations inside the store.

The idea of tracking users in the physical world might set alarm bells with regard to user privacy, or the lack thereof. While privacy

[*]Nandakumar (University of Washigton) was with MSR India as a research assistant, and Aggarwal (IIIT Delhi) and Goenka (IIT Kharagpur) as interns, during their participation in this work. Rallapalli (UT Austin) also started her involvement as an intern at MSR India.

is no doubt important, the experience in the online world shows that users are often willing to give up some privacy in return for free content and services, and even subsidized hardware and software (e.g., the Android OS). For example, when a user sends email regarding a disease, they might start seeing ads for treatments of that disease. So long as this information does not make its way to a party (e.g., a health insurance provider) who can have a material impact on the user, users seem willing to tolerate it. We argue that tracking in the physical world does not fundamentally alter this balance of privacy and economics. To mitigate user concerns, however, we advocate that tracking and analytics be performed only in public spaces, where users have to anyway contend with reduced privacy, because of the presence of other users, security cameras, etc. Users can opt out whenever they choose to.

Finally, the traversal of the physical world by a user, which points to the user's profile, clearly goes beyond the user's location. The user's history, gaze (e.g., what they are looking at), actions (e.g., picking up an item from a store shelf), physical context (e.g., light or sound level), social context (e.g., who else they are with), and even mood (e.g., whether they are happy or unhappy) all have a bearing on the true profile of a user. While we touch on some of these briefly, we focus on location, which we believe is the single most important dimension of a user's context, one that even sheds light on the some of the other dimensions listed above.

Our goal in this paper is to argue that Physical Analytics represents a new and exciting frontier for research on localization and other aspects of user context, providing a different and compelling context for such research and also broadening the goal of localization beyond just determining the user's coordinates. We present scenarios where Physical Analytics could provide value, and also discuss some specific research problems that we are pursuing.

2. OPPORTUNITIES

We now discuss various opportunities for Physical Analytics from the viewpoint of providing value to users and businesses. This is not intended to be an exhaustive list. Rather our goal here is to present some scenarios that support our position that physical analytics can be valuable, both on its own and in conjunction with online analytics.

In-store browsing: Retailing is big business, accounting for $15 trillion per year worldwide [14] Fundamentally, retailing is about connecting end consumers with the products and services in a way that serves the needs of both consumers (e.g., good value for money) and businesses (e.g., good sales). Clearly, therefore, deep knowledge of the consumer (or user) — what they are looking for, what competing offerings they have considered, their price sensitivity, etc. — is invaluable.

Analytics plays a significant role in online retailing, with clickstreams being mined, for instance, to make product recommendations. Such analytics draws on both purchases made by users and their browsing behavior, even if it ends short of an actual purchase. The situation is quite different in the case of physical, or brick-and-mortar, retailing, which still accounts for an overwhelming 96% of the global retailing market [14]. Physical retailers have little visibility customer behavior, even within the confines of their own stores. While customer loyalty programs provide some information, this is confined to the *purchasing* behavior of users; the considerable amount of time that users might spend *browsing* is lost in the blind spot. Retailers do not even have information on the average length of a store visit by a customer, and have to resort to low-tech and expensive means such as deputing sales representatives to tail customers as they browse through the store [24].

Therefore tracking and analyzing customer behavior in a physical store can be valuable. Indeed, this has become an active space for startups [16, 22] that are looking to leverage existing technologies for indoor localization (e.g., based on WiFi, video, LED light). The interesting research challenge, however, is in combining complementary sensing modalities, whether infrastructure-based or mobile-based, and performing meaningful analytics despite noisy data. For instance, can we perform accurate analytics in the aggregate despite errors in the underlying measurements (e.g., localization error)?

Dynamic pricing: Dynamic pricing is considered as the future of retail [15]. Based on a customer's behavior in a physical store (*e.g.,*whether and how long she browses clearance sections), the price can potentially be adjusted on the fly.

Space and event planning: Related to tracking in-store browsing is the question of how people flow through a space such as an airport, mall, or tradeshow floor. The owners of such spaces would be interested in answering this question so that they can optimize the layout for the future. For instance, the owners might wish to plan emergency evacuation plans based on where people tend to congregate, charge rent based on footfall concentrations, or locate facilities where more visitors are likely to benefit. In Section 4, we show an example of such analytics in the context of a demo event at Microsoft Research India.

Physical conversion: In advertising, conversion refers to the transition of a user who is shown an ad to a customer who makes a purchase or takes some other action indicative of an enhanced interest in the offering being advertised. Conversion tracking is the key to measuring the effectiveness of advertising campaigns. In the online world, conversion is tracked using HTTP cookies and tracking pixels to determine that a user went from being shown an ad to the checkout page on an e-commerce site. However, in many cases, conversion involves a physical action by the user, e.g., a visit to a car showroom or a doctor's office, or even more fine-grained level, a visit to the section of the Microsoft Store displaying the latest Surface tablet. So physical conversion tracking would be important. Indeed, businesses try to do such tracking today using survey forms, asking customers questions such as "please tell us how you heard about us." Obviating the need for such manual input holds the promise of making physical conversion tracking more accurate and complete. A mobile device that goes where the user goes could be the vehicle for tracking physical conversions. A key challenge, however, is linking online interactions made using one device (e.g., the user's home PC) with physical conversions tracked using a different device (e.g., the user's smartphone).

Better training data: Related to physical conversion tracking is the opportunity to augment training data with physical ground truth. Say we are looking to predict the commute habits users or whether they are interested in clothes, based on their online interactions. So, for instance, the presence of certain keywords in online searches by a user could be used as signals for such a prediction. However, we need labeled training data to fine-tune the prediction algorithm based on these signals. The labels, however, would have to come from the physical world, whether it is by tracking a user's commute or a visit to the clothing section of a store. Note that the labels thus obtained from a set of users would be used to fine-tune the algorithm to make better predictions in the future.

Personal assistant: Personal Assistant services such as Apple Siri and Google Now try to answer user queries and anticipate user needs in a manner that is natural and low-effort for users. For the same reason that Physical Analytics can benefit a business — by shedding light on user activity that is not visible online — it can also enable a personal assistant serve the user better. Physical information is already tapped by some personal assistants, for instance, to determine the locations of a user's home and workplace. However, there is the potential to gain deeper insight using fine-grained indoor location information. For example, if the user is found to have browsed through the lighting sections in multiple stores, they could be offered additional information or tips. They could also be connected (anonymously) with others users who share the same "physical trail", i.e., had spent time visiting the same sections in the recent past, and who thus may be in the position to answer any specific queries that the user might have.

3. RESEARCH CHALLENGES

Physical analytics could be enabled in two different ways, i) through assistance from infrastructure of interest (e.g., a retailer installing specialized equipment within the store such as video cameras) and ii) through assistance from the users' mobile devices (e.g. a mobile phone using its WiFi or inertial sensors). Having infrastructure support allows sophisticated instrumentation and mechanisms for potentially obtaining very accurate and fine grained analytics regarding users. For example, one could instrument shelves in a retail store to determine which items are being touched by the user. Having control or access to the users' device, on the other hand, allows for potentially obtaining information relevant to the user and even influencing user behavior. With increasing interest in wearable computing devices, such as Google glasses, the diverse ecosystem and capabilities of user owned mobile devices is bound to increase. Each of these kinds of assistance comes with its own research challenges.

Challenges in Infrastructure-Based Physical Analytics. There are two key research challenges to be addressed in infrastructure-based physical analytics – *Smart Infrastructure with Multi-modal Sensing* and *Privacy*.

Smart Infrastructure with Multi-Modal Sensing: The ability to directly control infrastructure gives a unique opportunity to innovate in terms of using multiple sensing modalities *e.g.,*WiFi and Video and creating innovative forms of smart infrastructure *e.g.,*smart shelves to improve the quality of information.

Privacy: Since Physical Analytics involves tracking user movement, it would likely raise privacy concerns. As in the online world, we believe that given an adequate incentive, users would be willing to put up with some loss of privacy. For instance, just as users swipe a loyalty card at the checkout counter today — gives up some privacy in return for a price discount — they could be incentivized to run a store app on their smartphone, say while in a store, so that their movements can be tracked more finely than would otherwise be possible.

Challenges in User-Based Physical Analytics. There are two key challenges to user-only based physical analytics – *Crowd Sourcing* and *Battery Power*.

Crowd Sourcing. Indoor location information is often of little use in the absence of context. For example, even knowing the lat-long of a user precisely will not reveal that the user was in fact in the shoes section. So there is the need to label indoor spaces in a semantically meaningful manner. Crowdsourcing can play a key role here, just as with outdoor maps (e.g., OpenStreetMap [5]). To limit the crowdsourcing effort, we would want to recognize places and detect movement automatically as users move about, and only prompt (some) users occasionally to obtain semantically meaningful labels for places. Such crowdsourcing throws up new challenges, e.g., the diversity in labels assigned by users for the same place (e.g., "shoes", "footwear", "sneakers"), and would require algorithms to clean and aggregate user-entered labels appropriately.

Battery Power To the extent that we are leaning on mobile devices to provide data for physical analytics, we need to be cognizant of battery life concerns. To a first degree of approximation, battery drain would be caused by both the use of sensors and by the processing of the sensor data. The latter can be streamlined by leveraging recent advances in heterogeneous multicore processing [1] and dedicated sensor processors [4, 3]. However, the energy cost of sensing require careful attention going beyond previous approaches towards triggered sensing (i.e., using a less expensive sensor to selectively trigger a more expensive one). We could, for instance, leverage the statistics of the patterns we are trying to mine themselves to reduce energy cost. For example, we might learn that users tends to dwell for an hour or two after entering a particular location (e.g., a movie hall), so the quality of physical analytics may not suffer even if we were to employ a high degree of temporal subsampling.

4. PHYSICAL ANALYTICS IN THE IDEAL CASE

In the ideal case, physical analytics would rely on both infrastructure support and user support in the form of a presence on the user's mobile device(s). Consider the problem of tracking customer movement through a retail store. The store has WiFi coverage, as is increasingly the case. A customer carries a smartphone running an app that is in a position to perform WiFi scans and tap other phone sensors. Such an app running on the customers' smartphones, coupled with information in the possession of the infrastructure owner (e.g., floor map, WiFi signal strength database or model), would enable localization of the customers' phones, thereby enabling tracking of the customers' movements through the store.

We show an example of analytics based in such tracking in Figure 1 in a demo event held at Microsoft Research India in Jan. 2013 with various demos spread across various locations on the floor, which measures 70m by 50m. From the perspective of physical analytics, there are a number of questions of interest, such as where people tend to congregate, how they tend to move between demos, etc. To answer these questions, we handed out phones loaded with a tracking app to 20 visitors and requested them carry the phones with them as they walked the demo floor. The app performed WiFi scans and then invoked the EZ service [13] in the backend to perform localization on a continual basis. The set of location estimates thus obtained across devices and time is then used to generate the heat maps are shown in Figure 1.

Figure 1a shows the raw heat map, and represents the density of user presence at each location over the entire duration, with the color going from light yellow (low density) to deep red (high density). Figure 1b shows the result of performing k-means clustering, with the cluster radius limited to 3m and clusters with too few points removed, thereby eliminating the outlier points. We were then able to label the clusters to indicate semantic information, i.e., the name of the demo or other activity associated with each cluster; the figure shows labels for the clusters corresponding to the Phytics demo and the food stand. We see that the foot stand cluster had the highest density of location reports, which raises the question of whether the food was a greater attraction than any of the demos! The number of location updates corresponding to each cluster depends both on the number of users that visited the cluster and the length of time they spent there. Figure 1c shows the heat map based on the average duration spent by a user within each cluster, which we treat as a proxy for the degree of user engagement. We see that the food stand cluster is no longer at the top — in fact, it is lighter than most other clusters — which indicates that the aver-

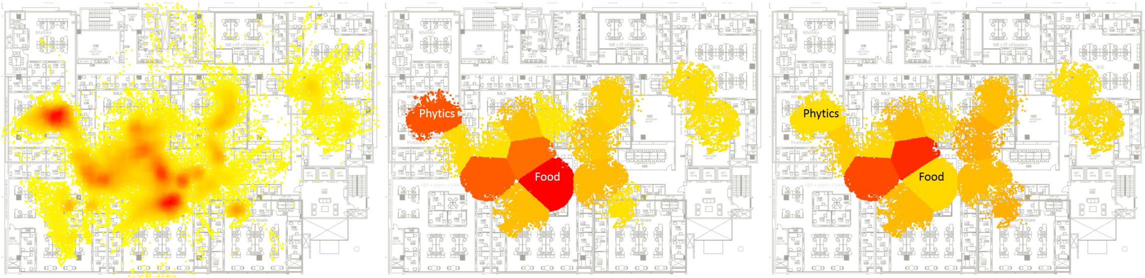

(a) Raw heat map (b) Heat map with clusters colored by location update count (c) Heat map with clusters colored by average time spent

Figure 1: Heat Map

age time spent there is low. Physical analytics reveals that the food stand cluster standing out in the overall heat map (Figure 1b) was because a large number of users visited that location; on the other hand, many users came to the food stand but they did not linger there for long, spending time at the demo locations instead.

As we can see, interesting physical analytics can be performed in an indoor setting based on information from both the user side (e.g., WiFi, accelerometer, and other sensor information from their mobile device) and infrastructure side (e.g., a WiFi infrastructure, WiFi signal strength model or database, floor map, semantically-meaningful location labels). Such a situation of cooperation between users and the business owning a space would be ideal and could be realized, for instance, by having a store provide incentives for customers to run a "loyalty card" app on their smartphones as they browse through the store, akin to the incentive for swiping a loyalty card at checkout time.

5. INFRASTRUCTURE-ONLY PHYSICAL ANALYTICS

Infrastructure-only Physical Analytics refers to physical analytics performed by the entity that controls the infrastructure in the physical space of interest, e.g., a retailer, a mall operator, or tradeshow organizer but with no direct cooperation from the the user's device. The lack of user cooperation eliminates a key source of information (e.g., the WiFi scans and accelerometer data in Section 4. So the key challenge is to find infrastructure-based substitutes for these and design algorithms to fill in for the missing information. In the rest of the section we provide two examples – *WiFi-Vision fusion* and *Smart Objects* to illustrate the challenges and opportunities unique to infrastructure-only physical analytics.

WiFi-Vision fusion: Given that user mobile phone's do not cooperate, the infrastructure could rely on sniffing passive background transmissions (e.g., 802.11 probe requests) from users' mobile phones. These sniffers could be arbitrarily sophisticated, for example they leverage recent advances such as using OFDM sub-channel level information [30] and antenna arrays [35] to obtain the best possible location estimates. However, relying only on background transmissions presents a significant challenge – sporadic nature of these background transmissions. For example, based on measurements made in a couple of malls in the Seattle area, we found that the median inter-arrival time between probe requests is 20-40 seconds, and the 85th percentile is over 3 minutes. In other words, the opportunity for the infrastructure to obtain a WiFi-based location fix on a mobile device will likely still be much less frequent than if the mobile device itself were scanning for beacons.

Vision-based tracking using cameras is a possibility, especially considering the prevalence of security cameras in indoor spaces. In principle, vision-based analysis can provide accurate tracks. However, our trials with several state-of-the-art commercial as well as open-source vision-based trackers revealed several challenges. The visual tracks tend to be fragmented, because of limited camera coverage, occlusion and even parts of a person's body being detected as distinct objects. Furthermore, there tend to be a large number of spurious objects detected, e.g., reflections off of glass surfaces and changing content on an electronic displays such as TV screens within view of the camera.

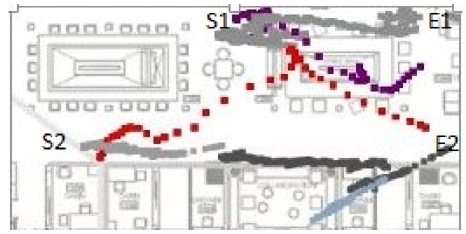

Figure 2: Combining vision tracking with WiFi tracking for 2 users, one who walks from S1 to E1 and the other from S2 to E2.

A combination of WiFi and Vision could provide best of both worlds. The continuous tracking capability using video can help fill up the sporadic nature of infrastructure based WiFi. Infrastructure based WiFi on the other hand can help fuse fragmented vision-tracks by establishing a consistent ID for each user (i.e., the wireless MAC address).

Figure 2 shows the results from a simple experiment, wherein two users, each carrying a phone, walked across a 30-meter long section of the floor, in a straight line from from left to right, one user from S1 to E1 and the other from S2 to E2. A camera at the right edge recorded the scene. The figure shows the result of processing both the WiFi and visual data. The purple and red squares show the WiFi tracks for the user 1 and user 2, respectively, each showing an error of meters. (In this experiment, we had frequent WiFi communication, so the WiFi localization was performed much more frequently than it would have otherwise been.) The streaks in grey correspond to visual tracks. We see that the tracks of each user get broken up into multiple "tracklets", shown in different shades of grey, because the same person is detected as a different object at various times. Furthermore, some of these tracklets extend into offices because of reflections off glass surfaces. Thus, we see that even this simple setting with just 2 users presents formidable challenges to leveraging vision-based tracking and the potential advantages of multi-modal techniques such as WiFi-Vision.

Smart Shelf and Smart Objects: The infrastructure could be instrumented to track more than just the locations of users, i.e., a richer set of interactions than just "footsteps". We briefly discuss two specific examples: smart shelf and smart objects.

The idea of a *smart shelf* is to track interactions of a user with objects on a store shelf. Inspired by our prior work on estimating road traffic density using cameras [28], we accomplish this using a pair (or, in general, an array) of inexpensive webcams mounted directly above the shelf, looking down on a virtual strip of flooring adjacent to the shelf. Whenever a user reaches out to touch or pickup an item, the user's arm would partially obstruct the cameras' views of the virtual strip. The location of the obstruction can then be used to calculate the location of the interaction on the shelf, which can then be coupled with information on the layout of items on the shelves. Our initial evaluation shows the promise of this approach, with both the location of an interaction on a shelf and the nature of the interaction (touching an item vs. picking it up) being detected accurately.

While a smart shelf helps detect basic interactions with "dumb" objects, *smart objects* can themselves be instrumented. Examples include powered computing objects, such as smartphones and tablets on display at an electronics retail store. An instrumented device can track a range of interactions with customers, e.g., how often the object was picked up, how long it was interacted with before being put back on the shelf, and even what functionality was involved while the customer was "kicking the tires".

With both dumb and smart objects, it is valuable to be able to tie interactions together, e.g., know that the user who just picked up a game title from a shelf is the same person who was checking out a gaming console a few minutes ago. User movement tracking would serve as the glue that binds together such disparate interactions.

6. USER-ONLY PHYSICAL ANALYTICS

For several indoor environments assistance or cooperation from the store may not be available. For example, the store may not allow placement of equipment such as cameras or provide access to feeds from them even if they are present. It may not provide any details of the product layout information or the floor plan. Further, most WiFi-based localization techniques depend on a training phase that require prior profiling of the area (*e.g.,*war driving or model building). In practice it is not practical to profile every store. In this section we ask the question *Can we perform physical analytics without the benefit of being able to locate the user at all and relying only on information from end users' mobile devices?* We answer the question in affirmative starting with one example – Label-Space localization.

Example : Locating in Label-Space: Consider a crowd-sourcing application where participating users are somehow incentivised (example: discounts) to enter *labels* relevant to their location. For example, suppose that a user enters a departmental store and occasionally stops and enters as *labels* such as "electronics section", "cameras", "perfumes", etc. Further, their phones are instrumented to collect WiFi scans from the phone in the background as they traverse the shop. The application is thus able to associate these *labels* with WiFi signal strength measurements. Further, if inertial sensor data were also collected in the background, it could be leveraged to provide a rough spatial distance estimate between two areas such as electronics section and perfumes section for example, by counting the steps taken by the user. Once there is enough data has been crowd-sourced from a certain store users can be located in this *label*-space by comparing their WiFi-fingerprint and steps taken. For example, by comparing WiFi-scans in the crowd-sourced database once can infer that the person is in the electronics section or the perfume section in a manner similar to location inference.

This manner of localization in the *label* space can provide very rich information for Physical analytics. For example, the user could be profiled based on the sections the he/she spends most of the time. The user could then be provided relevant discounts as soon as he/she approaches their *labels* of interest. Thus, in this manner physical analytics can be obtained without any support from the indoor space or localization capability.

A unique aspect of localization in *label* spaces is the inherent category hierarchy in *label* space. For example, the *labels* "camera", "T.V", "tablet" etc. all are encompassed within the category of "electronics". Similarly, the categories "SLR camera" "Macro Lenses" etc. are all encompassed within the "camera" category. Thus, when a user is localized in label space, he/she could belong to multiple categories along the category hierarchy tree. In fact, an ideal *label* space localization scheme would return the probabilities that the person belongs to each of the categories. Typically, the more general a category the larger the physical spatial extent of the category in the store and the higher the accuracy.

A label space localization experiment: As a first step towards enabling *label*-space localization we conducted experiments in three different indoor shopping areas – *Departmental store*, *Grocery store* and *Mall*. In each of these location was asked users to walk about carrying their mobile phones while briefly stopping at various locations to jot down relevant labels using an android app. The phone also continuously scanned WiFi signal strengths in the background. Labels could be associated with WiFi signals using time stamps. Data was collected at each of these locations on two different days, 3 days apart. Data collected on the first day was used as training data set while that on the second day was used to test.

Departmental Store. The store comprised two levels – with an area of about $30m \times 20m$ on level 1 and $30m \times 30m$ on level 2. The various categories explored in the store and the labels are shown in Figure 3 (a).

Grocery store. The store has only one level comprising an area of about $30m$-by-$40m$. The various categories explored in the store and the labels are shown in Figure 3 (b).

Mall. For these experiments we focused on small stalls, each situated around $7 - 8m$ apart from each other as depicted in Figure 4 (the layout was available in an online map). In these experiments we varied device type and carrying positions in order to evaluate the effects of such variations on performance. We experimented with two different mobile phones – Nexus 4 and Pantech Crossover and two different carrying positions – inside pant pocket and in hand.

Label Determination. In our testing, we tried to predict the *label* given the WiFi scans by user. In general this can be achieved by computing a distance metric between the WiFi scans in the training set corresponding to each *label* and the given WiFi-scans from the user. The *label* with the least distance is deemed the relevant label corresponding to the user in question. In our implementation we used Euclidean distance between the Received Signal Strengths (RSS) received from various APs as the distance metric[1].

Performance in Departmental Store. All 10 labels matched correctly in this store matched correctly thus matching at category as well as sub-category level showed 100% accuracy. The high accuracy was because there were two floors in the departmental store and the average distance between any two *label* areas in the same floor was about 15m. In addition to this large spread between categories, the average number of APs observed per WiFi-scan was quite large – about 14.

Performance in Grocery Store. There were 13 different *labels* collected from the grocery store and the overall accuracy in predicting the right *label* was about 76.9%. The accuracy in predict-

[1]Whenever measurements from certain APs were missing we used -100dBm as the RSS

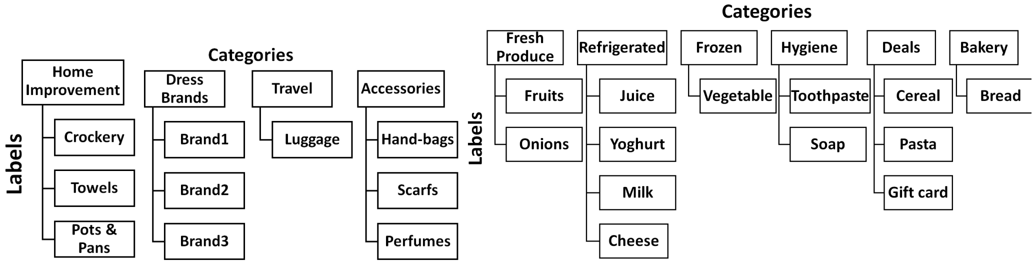

(a) Departmental store

(b) Grocery store

Figure 3: Categories

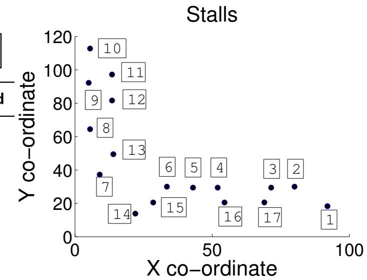

Figure 4: Spatial layout of dwell points

Experiment	Correct Match	Within Three Closest Stalls
Same device & placement	88%	100%
Same device, different placement	77%	94%
Different device & placement	64.8%	82%

Table 1: Fingerprint matching across walks

ing the right category, however was about 92.3% indicating that most items of the same category are usually adjacent to each other. The reduction in performance compared to the departmental store is because there is much more variety within the same space in a grocery store leading to different labels being spatially very close (an average of 8m). Further the average number of APs per scan was also significantly less compared to the departmental store at an average of 7.

Performance in Mall. We considered each small stall as a single *label* in this experiment. The performance results are depicted in Table 1. As seen from Table 1, when trained and tested for the same device and placement the accuracy of exact *label* is high however it deteriorates significantly when trained and tested over different devices and placements. Nevertheless, the *label* is matched correctly up to within two nearest stalls with very high accuracy despite device and placement variation. Such error is tolerable for several applications such as giving location specific discounts since, the correct stall is still probably very close if not the correct one.

7. RELATED WORK

Past work related to physical analytics can broadly be classified into two categories: startups, and research. A number of startups are building end-to-end physical analytics systems that focus on extremely niche business models. Past research, on the other hand, has focused on the individual components to physical analytics (e.g., localization, sensing, energy efficiency) rather than the complete end-to-end system.

Start-ups: Nearbuy Systems [22] requires retail stores to deploy their customized WiFi localization infrastructure for analytics on in-store customers (dwell times, visit frequency, window conversions, etc.) Euclid Analytics [16] leverages existing in-store WiFi infrastructure to provide similar analytics to retail stores. Mondelez [31] requires retail stores to deploy cameras in shelves that use facial-recognition to provide insights into demographics of customers that browse a given product. Because retail stores stand to benefit directly from these analytics, it is natural for these start-ups to assume extensive infrastructure support (densely deployed access points, detailed shelf maps, etc.) Their technologies do not generalize beyond the consumer-retail niche.

Localization and Sensing Research: There has been extensive work in indoor-localization and sensing [33, 34, 25, 17, 19, 7] that cover spaces both big (e.g., buildings) [32, 27] and small (e.g., retail stores) [12, 10, 30], both using infrastructure (e.g., WiFi, RFID) [9, 20, 23, 35] and not (e.g., purely-phone based, crowdsourcing) [8, 21, 13], sensing both the environment [26, 36] and users [29, 11]. There has also been recent work on combining video tracking with WiFi tracking [37], although the practical difficulties such as fragmented tracks and passive sniffing are not considered. With regard to analytics, there has been work on inferring traffic [2] and home/work locations [18] by using cellular information for wide-area tracking. Finally, there is ongoing work [6] building an urban livestyle innovation platform by collaborating with malls and recruiting university students as participants.

Physical analytics leverages past work in localization and sensing (when applicable) as components of the bigger system. Furthermore, because localization and sensing for physical analytics is not restricted to co-ordinate spaces and has lower accuracy constraints, physical analytics can leverage future localization systems (e.g., based on label-spaces) that target novel design points.

8. CONCLUSION

We believe Physical Analytics is a promising new direction for research on localization and other aspects of user context determination. Tracking user movement, particularly in indoor spaces, such as malls and stores, can reveal deep insights into a user's profile, just as clickstreams yield valuable information in an online setting. By articulating the opportunities for impact, the research challenges, and our initial work in this space, we hope to kindle broader interest in the research community in Physical Analytics.

Acknowledgements

With regard to combining vision-based tracking with WiFi-based tracking, we thank Shree Nayar for pointing us in this direction and Ratul Mahajan for making available his video tracker.

9. REFERENCES

[1] ARM biglittle. http://www.arm.com/files/downloads/big_LITTLE_Final_Final.pdf.

[2] Bangalore Transport Information System. http://www.btis.in/.

[3] iphone5s. http://www.apple.com/iphone-5s/design/.

[4] Motox. http://www.motorola.com/us/FLEXR1-1/Moto-X/FLEXR1.html.

[5] Opensteetmap. http://www.openstreetmap.org/.

[6] SMU LiveLabs. http://centres.smu.edu.sg/livelabs/.

[7] Time Domain. http://www.timedomain.com/.

[8] M. Azizyan, I. Constandache, and R. Roy Choudhury. SurroundSense: Mobile Phone Localization via Ambience Fingerprinting. In *MobiCom*, 2009.

[9] P. Bahl and V. N. Padmanabhan. RADAR: An Inbuilding RF-based User Location and Tracking System. In *INFOCOM*, 2000.

[10] K. Bauer, D. Mccoy, B. Greenstein, D. Grunwald, and D. Sicker. Using Wireless Physical Layer Information to Construct Implicit Identifiers. In *HotPETs*, 2008.

[11] S. Beauregard and H. Haas. Pedestrian Dead Reckoning: A Basis for Personal Positioning. In *WPNC*, 2006.

[12] V. Brik, S. Banerjee, M. Gruteser, and S. Oh. Wireless Device Identification with Radiometric Signatures. In *Mobicom*, 2008.

[13] K. Chintalapudi, A. P. Iyer, and V. N. Padmanabhan. Indoor Localization Without the Pain. In *Mobicom*, 2010.

[14] Cushman and Wakefield. Global perspective on retail: Online retailing, 2013. http://annualreview.cushwake.com/downloads/01_Global_Perspective_on_Retail.pdf.

[15] The future of retail is dynamic pricing. http://www.zdnet.com/the-future-of-retail-is-dynamic-pricing-so-why-cant-we-get-it-right-7000021444/.

[16] Euclid Analytics. Answers and insights for qsr, coffee, specialty retail, and large format stores, October 2013. http://euclidanalytics.com/.

[17] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. E. Kavraki. Practical Robust Localization over Large-Scale 802.11 Wireless Networks. In *MobiCom*, 2004.

[18] S. Isaacman, R. Becker, R. Cáceres, S. Kobourov, M. Martonosi, J. Rowland, and A. Varshavsky. Identifying important places in people's lives from cellular network data. In *Proceedings of the 9th international conference on Pervasive computing*, Pervasive'11, pages 133–151, Berlin, Heidelberg, 2011. Springer-Verlag.

[19] N. B. John, J. Heidemann, and D. Estrin. GPS-Less Low Cost Outdoor Localization For Very Small Devices. *IEEE Personal Communications Magazine*, 7:28–34, 2000.

[20] P. Krishnan, A. S. Krishnakumar, W. H. Ju, C. Mallows, and S. Ganu. A System for LEASE: Location Estimation Assisted by Stationery Emitters for Indoor RF Wireless Networks. In *Infocom*, 2004.

[21] H. Lim, L. Kung, J. Hou, and H. Luo. Zero-Configuration, Robust Indoor Localization: Theory and Experimentation. In *Infocom*, 2006.

[22] Nearbuy Systems. In-store mobile commerce, October 2013. http://www.nearbuysystems.com/.

[23] L. Ni, Y. Liu, C. Yiu, and A. Patil. LANDMARC: Indoor Location Sensing Using Active RFID. In *WINET*, 2004.

[24] Paco Underhill. Why we buy: The science of shopping. http://www.amazon.com/Why-We-Buy-Science-Shopping/dp/0739341928.

[25] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The Cricket Location-Support System. In *MobiCom*, 2000.

[26] T. Rappaport. *Wireless Communications Principles and Practice*. Prentice Hall, 2001.

[27] P. Robertson, M. Angermann, and B. Krach. Simultaneous Localization and Mapping for Pedestrians using only Foot-Mounted Inertial Sensors. In *Ubicomp*, 2009.

[28] R. Sen, A. Cross, A. Vashistha, V. N. Padmanabhan, E. Cutrell, and W. Thies. Accurate Speed and Density Measurement for Road Traffic in India. In *ACM DEV*, 2013.

[29] S. Sen, R. R. Choudhury, and S. Nelakuditi. SpinLoc: Spin once to know your location. In *ACM HotMobile*, 2012.

[30] S. Sen, B. Radunovic, R. R. Choudhury, and T. Minka. Precise Indoor Localization using PHY Layer Information. In *ACM HotNets*, 2011.

[31] L. Stampler. That chocolate chip cookie in the checkout aisle is watching you. *Time NewsFeed*, October 2013. http://newsfeed.time.com/2013/10/16/that-chocolate-chip-cookie-in-the-checkout-aisle-is-watching-you/.

[32] H. Wang, S. Sen, A. Elgohary, M. Farid, M. Youssef, and R. R. Choudhury. No Need to War-Drive: Unsupervised Indoor Localization. In *Mobisys*, 2012.

[33] R. Want and et al. The Active Badge Location System. *ACM Transactions on Information Systems*, Jan 1992.

[34] A. Ward, A. Jones, and A. Hopper. A New Location Technique for the Active Office. *IEEE Per. Comm.*, 4(5):42–47, 1997.

[35] J. Xiong and K. Jamieson. ArrayTrack: A Fine-Grained Indoor Location System. In *ACM HotMobile*, 2012.

[36] M. Youssef and A. Agrawala. The Horus WLAN Location Determination System. In *MobiSys*, 2005.

[37] B. Zhang, J. Teng, J. Zhu, X. Li, D. Xuan, and Y. F. Zheng. EV-Loc: Integrating Electronic and Visual Signals for Accurate Localization. In *ACM Mobihoc*, Jun 2012.