

Lightweight Material Detection for Placement-Aware Mobile Computing

Chris Harrison Scott E. Hudson

Human-Computer Interaction Institute, Carnegie Mellon University
5000 Forbes Ave., Pittsburgh, PA 15213
{chris.harrison, scott.hudson}@cs.cmu.edu

ABSTRACT

Numerous methods have been proposed that allow mobile devices to determine where they are located (e.g., home or office) and in some cases, predict what activity the user is currently engaged in (e.g., walking, sitting, or driving). While useful, this sensing currently only tells part of a much richer story. To allow devices to act most appropriately to the situation they are in, it would also be very helpful to know about their *placement* – for example whether they are sitting on a desk, hidden in a drawer, placed in a pocket, or held in one’s hand – as different device behaviors may be called for in each of these situations. In this paper, we describe a simple, small, and inexpensive multi-spectral optical sensor for identifying materials in proximity to a device. This information can be used in concert with e.g., location information, to estimate, for example, that the device is “sitting on the desk at home”, or “in the pocket at work”. This paper discusses several potential uses of this technology, as well as results from a two-part study, which indicates that this technique can detect placement at 94.4% accuracy with real-world placement sets.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces, Interaction styles. C5.3 [Computer Systems Implementation]: Microcomputers. - Portable devices (e.g., laptops, PDAs).

General terms: Design, Human Factors

Keywords: Sensors, material detection, mobile devices, cell phones, laptops, PDAs, placement detection, location, context-aware, situationally appropriate interaction.

INTRODUCTION

Mobile devices are transported into a wide variety of environments and circumstances. Presently, devices know little about the context in which they operate, and so the user must explicitly manage their behaviors to fit the current setting. However, if we hope to accommodate the ever-increasing number of mobile devices and settings in which they are employed, these devices must make more use of context and adapt automatically to changing conditions [3,4,10]. Unfortunately, automatically selecting the right behavior is challenging, and could benefit from the capture and analysis of multiple dimensions of context.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UIST’08, October 19–22, 2008, Monterey, California, USA.
Copyright 2008 ACM 978-1-59593-975-3/08/10...\$5.00.

Considerable prior work has concentrated on location detection [1,2,10]. If a device knows where it is located, it can modify its properties for that general environment. For example, a personal cell phone could be put into silent mode when at work. Location detection can now be easily performed outdoors using GPS, or indoors using wireless networks (e.g., 802.11x [11] or cellular signals [8], as well as short-range wireless methods, such as Bluetooth [5] and RFID [7]). Other contextual sensing efforts have investigated for example, light sensors to detect indoor locality [9]. Additionally, humidity, temperature, ambient noise, device orientation, and other properties can shed light on a device’s environment, and when combined with sensors like accelerometers, can drive activity recognition [6].

The sensor described in this note identifies placement by considering the multi-spectral absorptive and reflective properties of proximate materials. We believe this sensor is particularly promising for use in augmenting location information because 1) it provides information on space immediately surrounding the device (i.e., materials it is touching or nearly touching), 2) it requires no external infrastructure to operate and 3) the resulting data is available at the point-of-use: the mobile device.

With such a sensor, it is possible to distinguish semantically different *placements*, such as in the hand, on the desk, in a drawer, or in the backpack or purse. These examples would all likely be well within the circle of confusion for pure location sensing. We can take advantage of this newly available contextual dimension to make mobile devices more aware and situationally appropriate.

EXAMPLE USES

We now discuss how placement detection can be used to make two common classes of mobile devices, cell phones and laptop computers, act more appropriately. In the interest of brevity, we do not discuss how each example can be combined with location and/or activity recognition. However, it is easy to imagine how each example could be augmented to have different properties when the user is in different locations or engaged in different activities.

Cell Phone

Screens on cell phones rarely turn off entirely. Even when inactive, it is common for them to display the time, battery life, signal strength, and other information (some phones even have screensavers). However, this information is only useful when visible. During the (potentially long) periods of time when the screen is hidden, energy is unnecessarily wasted. With placement detection, the user could set which

materials are coupled with contexts in which the phone is visually accessible. This could allow the cell phone, for example, to turn off the display when in a pocket or purse, while keeping the screen active when resting on the office desk or kitchen table. Furthermore, in the placements where the phone is visually accessible, it might be worthwhile to display enhanced information (e.g., next appointment), and not, for example, switch to a low-information screensaver.

To avoid accidental input, many cell phones lock their keypads. This is important in the microenvironment of bag or pocket, which can be in motion and have other objects that can press against keys. This is in stark contrast to the relatively controlled environment of a tabletop. However, even on these essentially accident-free surfaces, cell phones still lock their keypads, requiring extra interaction when the user comes to use the device. Placement detection could be straightforwardly applied to alleviate this.

Placement detection can be especially valuable for alerts. For example, vibration works best when close to the user (e.g., pocket). However, if the phone is buried deep in a bag, the vibration alert is unnecessary and a drain on the battery. Instead, armed with the knowledge of its placement, the phone could employ a stronger than usual audio alert. Also, bags tend to increase the amount of time needed to locate and extract a device. Thus, the cell phone could decide to provide additional rings when in less-accessible locations (it could even be adaptive, learning over time how long it takes to be retrieved in different placements). Conversely, fewer rings and softer audio alert could be used when in the pocket, a highly accessible location.

The behavior of visual alerts can also be informed by placement detection. Many cell phones activate their displays during an incoming call, often activating an energy-greedy backlight to be more visually salient. This is useful when looking for a phone in a bag. However, we do not typically *look around* in our pockets, depreciating the visual alert (once the cell phone is pulled out of the pocket, and can no longer sense the pocket material, it can switch to an outside or in-hand mode with visual alerts enabled). Similarly, one can imagine an office context where an auditory alert would be inappropriate (e.g., office desk, conference table). In these situations, the cell phone could employ an extra strong visual stimulus.

Laptop

Placement detection can also be used to augment a laptop's behavior. One example application is screen dimming, which laptops employ in an effort to conserve power. Unfortunately, this frequently occurs when the laptop is being used, for example, when the user is reading or viewing a graphic, or returns after a temporary distraction. However, there are clearly some placements where it should be less aggressive with this feature, with the most notable example being the lap – a placement that almost ensures the laptop is situated in the user's field of view.

We can also use placement detection to change the laptop's security settings. For example, at home, a secure environment, the laptop can be set to never lock the screen (i.e.,

require a password to regain access). However, at work, where there may be strangers, the laptop could be set to hide the screen contents after five minutes of inactivity and require a password after ten. Even stricter settings could be associated with newly encountered materials, perhaps those at a coffee shop or conference center.

SENSOR

Our system relies on the fact that there is a particular set of materials associated with any placement. For example, a backpack might be made from a dark nylon, while a table at home is made of oak. These materials have widely varying spectral properties, which can be used to uniquely identify them. Indeed, humans often use color to identify materials, distinguishing, for example, the difference between pine and cherry wood.

To identify materials, we use two different light sensitive components in our prototype sensor (Figure 1). One is a common photoresistor, which varies its resistance in response to light intensity. The other is a TSL230 light-to-frequency converter, manufactured by Texas Advanced Optical Systems (see <http://taosinc.com>). This small, inexpensive IC offers highly accurate and programmable light intensity measurement in a small form factor.

Although sophisticated, these passive sensors alone are insufficient for reliable placement detection. In most cases, little to no ambient light penetrates the area where a device contacts a surface (e.g., the space between the bottom of a laptop and a table). This means that light intensity will be universally low for all placements. Thus, we must employ an active approach in our sensing.

Our sensor uses light emitting diodes (LEDs) to artificially illuminate the target material. The reflected light can be measured by our photosensitive components. This allows it to operate effectively, for example, in pockets, bags, and drawers, as well as on tables, bedside stands, and kitchen counters. We use five LED elements: infrared, red, green, blue, and ultraviolet. This enables us to accurately measure the reflective properties of a material in several distinct and useful swaths of the EM spectrum (including above and below what the human eye can perceive).

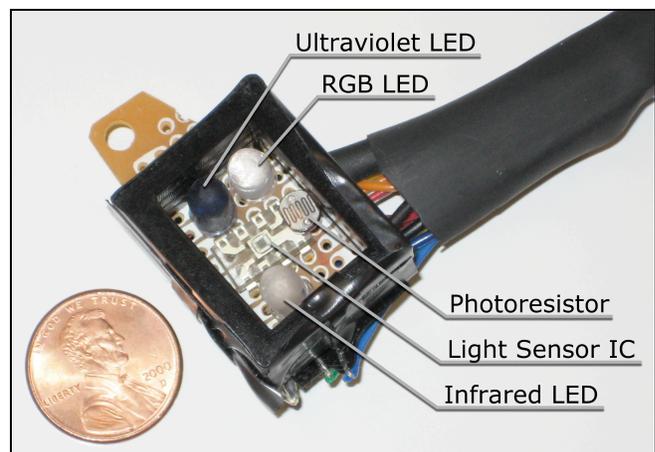


Figure 1. Prototype multispectral material sensor.

The sensor is controlled by an Atmel ATmega168 micro-controller, which operates as follows: The infrared LED is switched on and the reflective response from both sensing elements is recorded. Once this is complete, the LED is turned off. This process repeats for all of the remaining LEDs, each capturing data about a particular wavelength of light. Additionally, ambient light intensity is also recorded (i.e., no LED lit). This can be useful for sensing some materials (like a glass table), or rejecting measurements too badly affected by ambient light (perhaps allowing the device to wait for an accurate material measurement before assuming a behavior). The sensor has an effective range of approximately 2cm. Beyond this, ambient light tends to overwhelm the reflective response.

The entire sensing process can be achieved in under 25ms (50% of which operates outside the visible light spectrum – UV, IR, and ambient light readings). With sampling rates on the order of seconds, the technique is almost imperceptible to the naked eye, and thus is not visually distracting. Another benefit of this quick cycle time is low power consumption (~20mA when active). This could allow a device, for example, to do material detection once every 5 seconds or so with negligible impact on battery life. Furthermore, if mass-produced, this sensor could fit into a package smaller than half a cubic centimeter and cost less than a dollar, making it easy to incorporate into mobile devices. It may even be possible to take advantage of existing hardware on mobile devices, such as cameras and screens like those found on cell phones.

Figure 2 displays the spectral response of seven example materials. Data from both sensing elements are shown. It is apparent that the photoresistor (left) is far less sensitive to variations in lighting intensity than the more sophisticated light-to-frequency converter (right). Also, the particular photoresistor we used is insensitive to ultraviolet light (so only four values are shown).

It should be acknowledged that this type of multi-spectral sensor is not novel. Dozens of sensing approaches exist and numerous pre-packaged sensors are commercially available. They are most frequently used for simple color sensing applications, such as flame detection and currency identification. However, they are not typically used for general-purpose material detection because their accuracy is insufficient to support large material sets, like what might be required for inventory management at a clothing store.

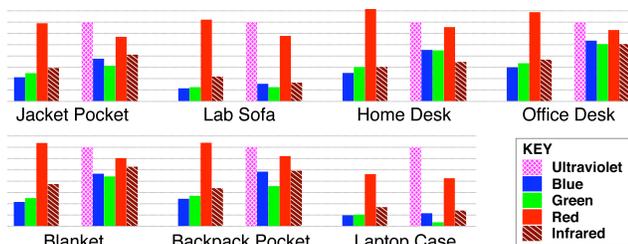


Figure 2. The multispectral data of seven materials. Each material is sampled by two independent sensors: a simple photoresistor (left) and a TSL230 light-to-frequency converter (right).

However, although poor for commercial-level spectrophotometry, they work well for identifying the limited number of places people put mobile devices - a context unique to HCI and ubiquitous computing.

We use a naive Bayes classifier for material identification. This application, built in Java and using the Weka machine learning toolkit (see <http://sourceforge.net/projects/weka>), reads data from the sensor in real-time over USB. The sensor provides eleven features: reflectivity of five wavelengths from the two sensing elements. The eleventh value is the ambient light intensity captured by the photoresistor.

This machine learning method requires training before it can be used. In a deployed scenario, this could be achieved by presenting a dialog to the user when encountering new materials. The user could then select which behaviors should be associated with that particular placement. However, because many materials are transient, this procedure should likely only be triggered after repeated encounters, ensuring that only routine placements are learned.

EVALUATION

To test the feasibility of using material identification for placement detection, we conducted two experiments. The first aimed to assess how reliable the sensor performed with a wide range of materials, and how well it scaled. The second experiment asked participants to list where they place their mobile devices. We then evaluated the accuracy of our approach using these real-world placement sets.

Accuracy and Scale

We sampled 27 materials twice a day for three days. Each time a material was sensed, 25 data points were recorded over a five second period (allowing noise inherent in the sensor to be captured). This produced 150 data points per material (6 sets of 25). No two material samples were collected consecutively, ensuring the sensor was coupled differently (e.g., orientation, pressure, location, ambient light interference). This data was combined into six trials, with each containing all 27 materials. We then trained our Naive Bayes classifier on five of the six trials, and used data from the remaining trial to evaluate the classification accuracy. This was done for all combinations of trials. Table 1 shows the combined results of this evaluation. Overall accuracy stands at 86.9% ($K=0.864$), with more than half of the materials achieving accuracies over 90%. This is an excellent result for such a large set of materials, many of which seemed almost identical to the naked eye. (Conventional 10-fold random holdout cross validation, which tends to be optimistic, yields an accuracy of 94.8%.)

Real World Placement Sets

Sixteen participants were recruited using an email campaign (8 female, mean age of 30). First, users completed an online survey asking them to list where they most frequently placed four classes of mobile devices: laptop, cell phone, audio/MP3 player, and PDA/organizer. The latter category had no responses. Each device category was followed by a question asking them to consider how frequently the device encountered places they did not list (i.e., non-routine placements). The results of this survey, seen in

PLACEMENT	ACCURACY	CONFUSED MATERIALS
Backpack	68.7%	Laptop Case (31.3%)
Backpack Pocket	82.7%	Jacket Outer Pocket (16%), Magazine Table (1.3%)
Black Plastic	93.3%	Carpet (6.7%)
Blanket	99.3%	Jeans Lap (0.7%)
Carpet	82.0%	Black Plastic (17.3%), Kitchen Counter (0.7%)
Home Chair	90.7%	Jeans Lap (9.3%)
Light Hand	98.7%	Jeans Lap (1.3%)
Lab Desk	62.7%	Khaki Pants (33.3%), Magazine Table (4.0%)
Dark Wood Table	84.7%	Jeans Lap (15.3%)
Home Desk	100.0%	-
Jacket Inner Pocket	84.7%	Lab Sofa (15.3%)
Jacket Outer Pocket	100.0%	-
Jeans Lap	63.3%	Suit Pocket (32.7%), Khaki Pants (4.0%)
Jeans Pocket	83.3%	Laptop Case (16.7%)
Kitchen Counter	80.0%	Carpet (18.0%), Laptop Case (2.0%)
Dinner Table	91.3%	Khaki Pants (8.0%), Home Desk (0.6%)
Electronics Fab. Desk	100.0%	-
Lab Sofa	100.0%	-
Lab Wood Table	98.7%	Leather (1.3%)
Laptop Case	66.7%	Jeans pocket (18.7%), Suit Pocket (18.7%)
Leather	83.3%	Home Desk (15.3%), Light Hand (1.3%)
Metal Cabinet	96.7%	Backpack Pocket (2.7%), Dinner Table (0.6%)
Khaki Pants	83.3%	Backpack Pocket (16.0%), Metal Cabinet (0.6%)
Magazine Table	85.3%	Lab Desk (12.7%),
Suit Pocket	82.0%	Jeans Lap (18.0%)
Dresser	92.7%	Electronics Fab. Desk (7.3%)
Dark Hand	92.7%	Jacket Outer Pocket (5.3%), Jeans Lap (1.3%), Home Desk (0.6%)

AVERAGE ACCURACY: 86.9%

Table 1. Accuracy results using 27 test placements.

Figure 3, show that users interact with a limited number of routine placements, averaging about five across all device classes. Additionally, participants believed that a majority of their devices rarely encountered new placements.

We then coordinated with ten participants to sample the placements they had listed in the survey. To capture as many materials as possible, we visited participants at work and/or at home. We sampled the material associated with each placement they had listed (and we had access to) three times. Like in the first experiment, each sampling was separated in time to ensure a different sensor coupling. We trained our naive Bayes classifier on two of three samples, and tested on the remainder (all combinations). The results are shown in Table 2. We achieved an overall placement detection accuracy of 94.4% with an average of 4.3 placements (out of an average of 5.5 total placements). This performance is especially impressive considering the classifier trained on only two sample sets per placement.

CONCLUSION

We believe that placement detection offers a valuable contextual dimension that can compliment both location and activity recognition. We present a multispectral sensor used for material identification, from which we can often infer placement. Results from our study indicate such a sensor is surprisingly accurate, and can be used to inexpensively augment mobile devices.

ACKNOWLEDGEMENTS

This work was supported in part by grants from the Intel Research Council, General Motors, and the National Science Foundation under Grant IIS-0713509.

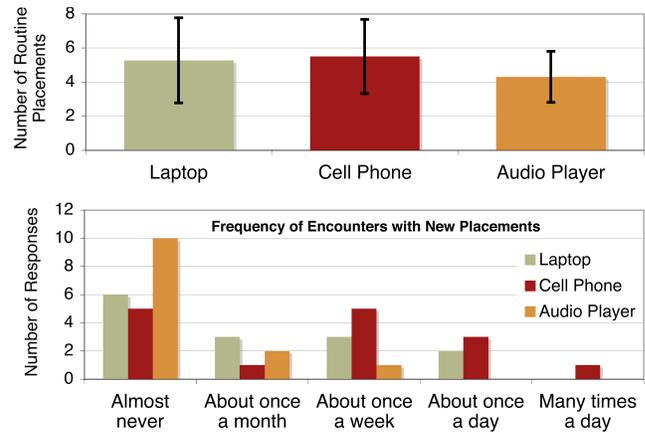


Figure 3. Number of routine placements (top) and frequency of encounters with new placements (bottom) for three classes of mobile devices.

REFERENCES

- Borriello, G., Chalmers, M., LaMarca, A., and Nixon, P. Delivering real-world ubiquitous location systems. *Communications of the ACM*, 48, 3 (Mar. 2005), 36-41.
- Hightower, J. and Borriello, G., Location Systems for Ubiquitous Computing, *IEEE Computer*, August 2001.
- Hinckley, K. and Horvitz, E. 2001. Toward more sensitive mobile phones. In *Proc. of UIST '01*, pp. 191-192.
- Ho, J. and Intille, S. S. Using context-aware computing to reduce the perceived burden of interruptions from mobile devices. In *Proc. of CHI '05*, pp. 909-918.
- Kawsar, F., Fujinami, K., and Nakajima, T. A light-weight indoor location model for sentient artefacts using sentient artefacts. In *Proc of SAC '07*, pp. 1624-1631.
- Maurer, U., Smailagic, A., Siewiorek, D. P., and Deisher, M. 2006. Activity Recognition and Monitoring Using Multiple Sensors on Different Body Positions. In *Proceedings of BSN '06*, pp. 113-116.
- Ni, L. M. et al. LANDMARC: Indoor Location Sensing Using Active RFID. In *Proc. of PerCom'03*, pp 407-415.
- Otsason, V. Varshavsky, A., LaMarca, A., Lara, E. Accurate GSM Indoor Localization, In *Proceedings of UbiComp '05*, pp. 141-158.
- Randall, J., Amft, O., Bohn, J., and Burri, M. LuxTrace: indoor positioning using building illumination. *Personal Ubiquitous Computing*, 11, 6 (Aug. 2007), 417-428.
- Schmidt, A., Beigl, M., and Hans-W, H. There is more to context than location. *Computers and Graphics*, 23, 6 (1999), 893-901.
- Wang, Y., Jia, X., and Lee, H. K.. An indoors wireless positioning system based on wireless local area network infrastructure. In *Proceedings of SatNav'03*, July 2003.

PARTICIPANT ID DEVICE TYPE	1	2	2	2	3	4	4	5	6	7	8	9	10	AVERAGE ACCURACY
	Cell	Laptop	Cell	Audio	Cell	Cell	Laptop	Cell	Cell	Cell	Cell	Cell	Laptop	
NUMBER OF ROUTINE PLACEMENTS	6	4	8	3	4	4	10	4	5	5	7	7	4	
NUMBER OF PLACEMENTS SAMPLED	4	4	8	2	4	3	8	4	3	5	3	5	3	
PLACEMENT DETECTION ACCURACY	97.7%	100%	75.2%	95.2%	75.0%	100%	100%	100%	99.1%	90.6%	100%	93.9%	100%	94.4%

Table 2. Accuracy results using participants' real-world placement sets.