From Position to Place

Jeffrey Hightower
University of Washington
Dept. of Computer Science & Engineering
Box 352350
Seattle, WA 98195 USA
+1 206 543 1695
jeffro@cs.washington.edu

ABSTRACT
Emerging proactive applications want to reason about "place", not coordinates. Existing systems rely on manually defining places which, while useful, does not scale to ubiquitous deployment. In this paper I define place and challenge the research community with learning and labeling places automatically.

Keywords
sensor fusion, location sensing, activity inference

INTRODUCTION
Several researchers including myself have created systems to fuse live measurements from multiple location technologies. Such systems provide a technology-independent location interface and allow probabilistic queries for objects' geometric positions and relationships. Applications like moving-map navigation and distance-aware buddy lists can easily be built on such infrastructure. However, emerging applications require a more symbolic notion: place. Generically, place is a human-readable labeling of positions. A more rigorous definition is an evolving set of both communal and personal labels for potentially overlapping geometric volumes. An object contained in a volume is reported to be in that place. This programming model maps well to event-driven application programming and is used by most existing location-aware computing frameworks including MSR Easy Living [1], AT&T Sentient Computing [2], and my own Location Stack work [3]. The latter also provides probabilistic confidence values for each labeling. One might argue to skip the geometric intermediary and determine place directly from sensor hardware, the common indoor example being infrared basestations corresponding to rooms, however this is unsatisfying due to the extensive engineering and rigidity of such an approach and has largely been rejected by the community as a general solution.

CHALLENGES
Current approaches require manual definition of places. I must, by hand, delineate and label my neighborhood, property, rooms, furniture, and service areas of my devices, for example, the area in which each of my display screens is visible. Then I can add specific semantics to build applications. Manual definition does not scale. Instead, ubiquitous deployment requires automatically learning significant regions and semantically labeling them as places. I pose these two challenges to the location research community and discuss work in progress.

Predicting Places from Maps and Behavior
The world has static structure such as roads, parks, rooms, and buildings. Maps capture this information well. The world also has dynamic physical constraints observable indirectly through the behavior of people and other entities. For example, people congregate in certain places at certain times, there are travel congestion points, and certain paths and are commonly taken from A to B. The challenge is to augment maps of physical features with the dynamic data to, over time, suggest geometric regions which are good candidates to label as places.

Promising work in this area uses automatic integration of maps, geographic information systems (GIS) data, and usage logs. By looking over time at where I go, how long I spend there, who else is around, and other things I do while there, my system can learn my hubs of activity and methods of transportation. The work in [4] applies this approach to learning typical modes and routes of transportation around a metropolitan area. In indoor environments, [5] shows how to compute the graph-like structure of rooms and hallways from maps to improve the performance of location estimation and enable path prediction. Also, the robotics community has explored the problem of automatically dissecting grid-based maps through robot exploration to learn detailed features and topological layout [6,7]. These efforts must continue with increasing emphasis on wide-area deployment and larger numbers of users.

Labeling Places
Labeling a geometric region assigns it semantics and can help prune and improve place predictions. More importantly, a label directly represents the place's demographic, environmental, historic, personal, or commercial significance and is the desired abstraction for emerging proactive applications. Manually labeling places does not scale so the research challenge is to automate the process.

Simple automated place labeling is already commercialized. Merging web data such as postal addresses with maps enables Nearest-X services where I can map and route myself to "coffee places" offering nearby purchase of espresso beverages – perhaps of a specific brand. Nearest-X services are useful, but they assume the learning input is mostly static and there is only a single type of place. New research seeks to relax both assumptions.

Labeling Places by Inferring Activities

This research seeks to automatically label places by inferring people's activities in those places. For example, here is the library where people select and read books, there is the kitchen where they cook food and wash dishes, and over there is an office where they use the computer and telephone. The Guide project at Intel Research Seattle is research in this space. Guide infers users' activities from knowledge of recently touched physical objects which are tagged with cheap RFID stickers and sensed by a wrist watch short-range reader. Guide mines the web offline via Google and eHow to build activity models and object-activity correlation probabilities, then can predict activities in real-time using a dynamic Bayes net sampling technique [8].

Guide is interesting because it also consumes place information and may be considered an application of place. The basic `isTouching(object,time)` primitive is easily augmented with `isAt(place,time)` and the web mining extended to populate place-activity correlation probabilities. Considering Guide an application is valuable because it highlights an important requirement for the place programming interface: The interface should be capable of answering both "What place labels are associated with my current coordinates?" and "What is the probability I am currently is in a place P?"

Labeling Places using Grassroots Contributors

Another research path starts by observing that the aggregate of many people periodically labeling their positions is a global place database. For example, if I occasionally provide a name quick or description of my coordinates, over time we can learn significant places by aggregating my labels with those contributed by other people. There are 3 challenges to this approach.

1. **Collaborative Filtering**. Like collaborative filtering for eCommerce web sites, a user-contributor place database must be robust to both incongruous and manipulative contributions. Meeting this challenge requires new research into combining reputation management with machine learning and location sensing. For example, much like Google's PageRank algorithm foils attempts to artificially inflate a given web page's search results, a place database must inherently resist similar attempts at manipulation by commercial or other interests.
2. **Data Management**. Making a scalable grassroots data management service for place information can call on the substantial expertise in the systems and databases communities around distributed peer-to-peer data management. It also requires developing the schemas and ontology of place data. For example, knowing that "diner", "restaurant", and "International House of Pancakes" are comparable entities allows refinement of the region in question.
3. **Human Interface**. Grassroots place contributor research must pay careful attention to interaction issues such as: What are the incentives for contributing to the communal database? How simple and transparent is the process of adding a place name? How are privacy concerns allayed?

CONCLUSION

Automatically predicting and labeling places is important because manual methods do not scale. Success in this research will enable more ubiquitous deployment of location technology and pave the way for revolutionary new applications which can reason about place instead of coordinates. To inform the work, we must create more applications which use place as a primary input for higher level inference, and, to evaluate scalability, it is critical these applications be widely deployed and have value to real users outside the research lab. The interdisciplinary PlaceLab program may be an ideal venue for this effort to move from position to place. PlaceLab is a grassroots effort to create a privacy-observant, planetary, indoor & outdoor positioning system with low barriers to participation. See `www.placelab.org and` [9] for more information.

ACKNOWLEDGEMENTS

Gratitude goes to Matthai Philipose, Gaetano Borriello, and Anthony LaMarca for vetting ideas and reviewing text.

BIOGRAPHY

Jeffrey Hightower is a doctoral candidate at the University of Washington in Seattle. His research interests are in employing devices, services, sensors, and interfaces so computing can calmly fade into the background of daily life. Specifically, he investigates abstractions and statistical sensor fusion techniques for location sensing. He received an MS in Computer Science & Engineering from the University of Washington in 2000 and is a member of the ACM and the IEEE.

REFERENCES

1. Brumitt, B., Krumm J., Meyers, B., and Shafer, S. Ubiquitous Computing and the Role of Geometry. *IEEE Personal Communications*, October 2000, pp. 41-43.
2. Addlesee, M., Curwen, R., Hodges, S., Newman, J., Steggles, P., Ward A., and Hopper A. Implementing a Sentient Computing System. *Computer* 34(8), pp.50-56, August, 2001.

3. Hightower, J., Fox, D., and Borriello, G. The Location Stack. *IRS-TR 03-008*, Intel Research Seattle, 1100 NE 45th Street, Seattle, WA 98105, June 2003.
4. Patterson, D. J., Liao, L., Fox, D., Kautz, H. Inferring High-Level Behavior from Low-Level Sensors. *Proceedings of the Fifth International Conference on Ubiquitous Computing,* October, 2003
5. Liao, L., Fox, D., Hightower, J., Kautz, H., and Schulz, D. Voronoi Tracking: Location Estimation using Sparse and Noisy Sensor Data. *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, IEEE/RSJ, 2003.
6. Schröter, D., Beetz, M., and Gutmann, J-S. RG Mapping: Learning Compact and Structured 2D Line Maps of Indoor Environments. *IEEE Intternational Workshop on Robot and Human Interactive Communication,* September, 2002.
7. Thrun, S. and Bücken, A. Integrating Grid-Based and Topological Maps for Mobile Robot Navigation. *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI)*, Portland, Oregon. 1996.
8. Philipose, M., Fishkin, K.P., Fox, D., Kautz, H., Patterson, D., and Perkowitz, M. Guide: Towards Understanding Daily Life via Auto-Identification and Statistical Analysis. *The 2nd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications*. Seattle, Washington, USA, October 2003
9. Schilit, W.N., LaMarca, A., Borriello, G., Griswold, W., McDonald, D., Lazowska, E., Balachandran, A., Hong, J., and Iverson, V. Ubiquitous Location-Aware Computing and the "Place Lab" Initiative. *The First ACM International Workshop on Wireless Mobile Applications and Services on WLAN*, September 19, 2003, San Diego, CA, USA.