A Quantitative Method for Revealing and Comparing Places in the Home

Ryan Aipperspach[1,2], Tye Rattenbury[1], Allison Woodruff[2], and John Canny[1]

[1] Berkeley Institute of Design, Computer Science Division,
University of California, Berkeley, USA
{ryanaip, rattenbt, jfc}@cs.berkeley.edu
[2] Intel Research Berkeley, USA
woodruff@acm.org

Abstract. Increasing availability of sensor-based location traces for individuals, combined with the goal of better understanding user context, has resulted in a recent emphasis on algorithms for automatically extracting users' significant places from location data. Place-finding can be characterized by two sub-problems, (1) finding significant locations, and (2) assigning semantic labels to those locations (the problem of "moving from location to place") [8]. Existing algorithms focus on the first sub-problem and on finding city-level locations. We use a principled approach in adapting Gaussian Mixture Models (GMMs) to provide a first solution for finding significant places within the home, based on the first set of long-term, precise location data collected from several homes. We also present a novel metric for quantifying the similarity between places, which has the potential to assign semantic labels to places by comparing them to a library of known places. We discuss several implications of these new techniques for the design of Ubicomp systems.

1 Introduction

The importance of understanding users' contexts is widely accepted in the Ubicomp community. Understanding context can help systems behave appropriately in a variety of situations and integrate more seamlessly into everyday life. Recently, the increased availability of position sensing technologies has resulted in a focus on finding significant places based on traces of a user's position [3, 8–10, 12, 13, 23]. Places are stable contexts within which social practices are situated [7], playing a significant role in broader definitions of context.

Hightower et al. [8] list two problems in developing algorithms for finding significant places for individual users: (1) finding significant locations and (2) assigning semantics or names to those significant locations ("moving from location to place"). The primary contributions of this work are a solution to the first problem in the domain of the home and the introduction of a novel method for addressing the second problem, based on a measure for comparing the similarity between places. Most existing algorithms for finding significant locations target building-sized locations within city-scale areas, since most available location data sets exist at a building-scale resolution. We have collected the first long-term, high-precision home location traces [2], enabling us to begin developing place-finding algorithms for the home. The distribution of places within the home differs markedly from the distribution at city-scale, motivating the development of new techniques for finding places in the home. To validate our algorithm choice, we discuss the places it finds in several homes.

Most existing place-finding algorithms focus on the problem of finding significant locations rather than on assigning semantic labels to the places they find. Instead, they rely on *geocoding* repositories (e.g. Microsoft MapPoint) or on hand-labeling by people [26] to apply labels. Hightower et al. [8] also note that geocoded information, like raw coordinates, "does not correspond to someone's mental model of their personal routine nor to the terminology they use when discussing the places they go." We provide a step toward solving this problem.

Concretely, we describe a method for measuring the similarity between different places that should enable automatic labeling of new places. By measuring the similarity between places, it is possible to find recurring structures that reveal the shared types of place within a particular culture (e.g., bedrooms in one home should be similar to bedrooms in other homes within the same culture). While it is difficult to learn the specific details of a place, we claim that the shared attributes of places can be automatically detected. Using these attributes, new places can be compared to a library of known places resulting in an (indirect) understanding of the new place. The limited availability of location traces from homes makes it difficult to build and test a library of similar places. In this work, we make use of the largest high-precision home location data set currently available to highlight the potential of such techniques. Our similarity metric is based on emergent properties of places resulting from people's patterns of behavior. Consequently, we claim that it satisfies some of the notions of context constructed through actions, rather than imposed via a fixed structure.

In section 2 we describe existing definitions of place and argue for the importance of measuring the similarity between different places. In section 3 we discuss related work. Section 4 describes our algorithm and the results of applying it to a set of location data collected in several homes. Finally, section 5 describes potential design implications and section 6 concludes and provides directions for future work.

2 A Definition of Place

2.1 Existing Definitions of Place

The definition of place that we use in this paper is most aligned with the work of Harrison and Dourish [7]. They observe that designers of interactive systems and virtual environments often build notions of "space" – the three dimensional environment in which objects interact and events occur – into their systems with the goal of facilitating interaction by framing users' behavior in the same way that it is framed in the physical world (i.e., by relying on spatial metaphors). However, they argue that behavior in the physical world is framed not only by space but also by the actions normally performed in that space.

Harrison and Dourish define place as "a space which is *invested with understandings* of behavioral appropriateness, cultural expectations, and so forth," saying that "we are *located* in space, but we *act* in place". That is, a place is useful for framing behavior because a user can make use of culturally embedded knowledge to determine what activities are appropriate for particular places. Whether walking into a church, a post office, a living room, or a kitchen, people use cues based both on both the architecture and configuration of the space and on the observable behaviors of others to understand what kind of place they are entering and how they should act. According to this definition, a place's meaning is jointly constructed by the physical structure of a location and the activities that regularly occur there.

This definition of place is a useful definition for designers of Ubicomp systems. Many recognize the need for determining users' context to better support their activities through appropriate and relevant actions [4]. They look for cues that determine appropriate behavior in particular situations, which is exactly what Harrison and Dourish suggest that "place" encodes. Interestingly, Harrison and Dourish's definition of place is similar to definitions in other fields, such as *place settings* in environmental society [27]. However, a discussion of these similarities is beyond the scope of this paper.

2.2 Working Definition of Place

From the existing definitions of place, we highlight six attributes of place along two different dimensions:

- *Attributes relating to time-invariant physical configuration*
  - **Position**, a bounded region in three-dimensional, physical space
  - **Physical structure**, the architectural components of a space
  - **Object co-presence**, the consistently present artifacts in a space (e.g. furniture)
- *Attributes relating to the historical behavior within a place*
  - **Time**, the distribution of time spent in a particular place
  - **Stable patterns of behavior**, the set of activities that occur in a place
  - **Person co-presence**, the patterns of presence of other people within a place

Drawing on activity theory [11], we argue that common patterns of interaction between people, objects, and locations (as captured in the six attributes above) embody the set of common behaviors characteristic to each place as suggested by Harrison and Dourish. The categorization of these attributes into two dimensions is a matter of conceptual convenience. In general, the physical configuration attributes and the historical behavior attributes are deeply intertwined. (Notions like "affordance" attempt to capture this connection.) In fact, their interconnection can be considered both a consequence of and a generator of social culture [7]. Practically, however, the distinction between physical and behavioral attributes prescribes a basis for the algorithms we present later.

Shared Notions of Place Because places are intertwined with culture, there exist common notions of place within specific cultures (e.g., "kitchen", "meeting room", "bedroom"). This fact is captured in the following argument: because strangers can enter spaces and, at a glance, assess both the names of the places and the activities likely to occur there, there must exist a shared notion of types of places and the appropriate actions for each of those types. (Of course, strangers may miss more idiosyncratic activities, e.g. eating cereal for breakfast while standing in the kitchen).

As a result of these shared cultural definitions of place, it should be possible to measure the similarity between particular "instantiations" of those places in individual homes – e.g., bedrooms should appear similar to other bedrooms, and kitchens should appear similar to other kitchens. In this work, we present an initial metric for comparing places in the home. We utilize the fact that different types of places tend to support different typical activities, basing our comparison metric on the patterns of use in the spaces we compare. The correlation between types of place and typical activity patterns is supported by previous work. For example, in a study of 100 Brazilian homes, Monteiro [15] found a correlation between the types of activities carried out in different places and the relative distance of each place from the entrance of the home. In section 5 we discuss several implications of a system for measuring shared notions of place.

3 Previous Work

Recently, a number of authors within the ubiquitous computing community have considered the problem of finding significant locations based on a person's location history. We describe some of the existing methods below. Most algorithms focus on finding significant locations at the city-scale. In agreement with previous studies of place within the home [18, 28], we argue that places in the home are significantly different from larger-scale places. Consequently, different methods are required for finding places in the home.

3.1 Existing Place-Finding Algorithms

Existing place-finding algorithms can be divided into two classes, *geometry-based* and *fingerprint-based* [8]. Examples of geometry-based algorithms include those by Marmasse et al. [13], Ashbrook and Starner [3], Liao et al. [12], and Kang et al. [9]. Each of these algorithms takes a history of locations from a single person (e.g. from GPS readers) and finds locations where the person spends significant periods of time. The algorithms vary based on the type of sensor data they use and on the specific clustering algorithm they use. For example, Ashbrook and Starner search GPS data for locations where position changed slowly or readings dropped out completely, indicating times when the user was not moving or was inside a building. They apply a hierarchical variant of k-means to cluster the candidate locations. Kang et al. use sensor-agnostic location traces to determine when users remained within a (specified radius) "space-sized" region for at least some minimum amount of time to find significant locations.

Examples of fingerprint-based algorithms include those by Trevisani and Vitaletti [23], Laasonen et al. [10], and Hightower et al. [8]. These algorithms calculate the "fingerprint" of a person's location (e.g., a vector of currently visible cell towers and wireless access points), find recurring fingerprints where the person spends a significant amount of time, and return a list of places based on significant fingerprints. These algorithms make the assumption that there exists a stable mapping between the fingerprints they find and the three-dimensional physical space that humans inhabit.

Existing place-finding algorithms find significant locations for different individuals, but they do not attempt to label types of places or to compare places across different people. For example, existing algorithms do not address the problem of automatically labeling places as belonging to the class "work" or "home" or of determining that two people's "work" places are similar unless those places are in the same physical location. Without knowledge of these distinctions, it is more appropriate to treat these methods as finding significant locations, not places.

3.2 Place in the Home

There has been little work done in determining significant places in the home based on location histories, primarily because high-resolution (sub-meter accuracy) location traces for residents within the home have not been available. However, several observational studies have explored the concept of place in the home. Oswald et al. [18] found that elders tend to spend much of their time in *favored places*. These centralized locations, which often provide easy access to essential (e.g., medication and food) and important objects (e.g., pictures and diaries) are crucial in allowing elders to remain autonomous as moving about their homes becomes more difficult. Based on qualitative studies in a range of homes, we found similar use of favored places within a more general population (see [28]). Like Oswald, we found that study participants spent the majority of their time in favored places. However, we found that participants also visited a range of *kinetic places*, or places used for shorter duration activities often involving physical manipulations, such as the bathroom, a mirror for doing one's hair in the morning, or a kitchen counter used to make sandwiches for lunch.

The work cited above suggests that places in the home exist at a variety of levels, ranging from small places like a single mirror to entire rooms. In this work, we are interested in finding places that are large enough to capture generic usage patterns across people and homes (e.g. cooking, sleeping), yet small enough to reveal differences in these patterns. Places satisfying these constraints should be comparable between people and be semantically meaningful. Often, different sized places exist even at the same semantic level (e.g., living rooms and bathrooms). The range of place sizes makes it difficult for existing algorithms that assume places at a single level have similar sizes (e.g. [9, 13]).

Additionally, places in the home are much closer together in relation to their size than are places at the building scale (consider the distance between work and home, compared to the difference between two rooms in a home). The small size and close proximity of places in the home means that the separation between places is close to the scale of errors in location sensor readings, even when using very accurate sensors. (The size of places and typical error bounds for our sensor data are both on the order of meters.) As a result, a more probabilistic view of places is necessary than is taken in many existing place-finding algorithms. For example, we applied Kang et al.'s place-finding algorithm to the data for each participant (with a place-size of 2 meters and a time window of 2 minutes). Because place sizes and noise in the home are at similar scales, readings taken while a user remains in a single place frequently fall outside the bounds of that place. Of data while users where home and awake, the Kang algorithm labeled an average of 45% of points as "moving between places" and 55% as belonging to a specific place. However, qualitative data analysis and interviews based on our data found that most people spend the majority of their time in a small number of favored places and little time moving around [2], suggesting this large portion of unlabeled data is an undesired artifact of the algorithm.

Consequently, we developed a new geometry-based algorithm for finding significant places in the home. This algorithm must work for environments where the ratio of place size to inter-place distance is high. Therefore, the algorithm must be capable of finding places without relying on the detection, either explicitly or implicitly, of movement paths between places.

4 Finding Place in the Home

As discussed earlier, we treat place as an emergent quality of the physical structure of a location and of the routine activities that people carry out there. Instead of directly processing the rich details intrinsic to both of these dimensions, we use simple proxies. For physical structure we use the set of locations that people have occupied. For action and activity we use the first-order temporal patterns in each person's position data. Both proxies are common input features for existing place-finding algorithms (e.g. [8, 9]), and the tunable parameters used by our algorithm are similar to those used by other algorithms. However, while these algorithms are reliable and accurate in their intended domains (city-scale areas), they do not work at the scale of places within a home (see section 3.2).

Like other place-finding algorithms, our main algorithm described in section 4.2 identifies significant locations rather than places. However, the similarity metric described in section 4.3 considers the activity patterns that differentiate places from locations, helping to address the problem of moving from location to place. While the location data in our data set does not capture the full range of attributes of place described in section 2.2, the spatial and temporal patterns provide a good starting place for such analysis. As richer data sets become available, the algorithm can be expanded to include additional attributes of place.

4.1 Data set

There are currently few data sets providing accurate in-home location information. Typical examples are Rowan and Mynatt [20], who collected one year of data from a single woman using in-floor pressure sensors, providing meter-level accurate location traces, and Tapia et al. [22], who collected 2 weeks of data from 2 single-person apartments using ∼80 binary state-change sensors with known locations, providing sub-room level accuracy.

In this study, we make use of highly-precise (sub-meter) sensor data which we collected from 3 homes using ultra wideband (UWB) sensors from Ubisense, Ltd. [25]. To our knowledge, this is the first long-term, high-precision set of location data collected in the home. The data set includes location traces for each person (7 total) in the study, over a time period ranging from three weeks to several months in each home. Each data point is a 4-tuple containing: (1) x position, (2) y position, (3) time, and (4) the duration of time that the measurement covers. For more information about the data collection process, see [2]. The participants from each home are described below:

1. *Home 1: Brad and Jacqueline.* A one-bedroom apartment occupied by Brad and Jacqueline, two graduate students from Australia.
2. *Home 2: Jack and Margaret.* A one-bedroom apartment occupied by Jack and Margaret, a recently married couple from England.
3. *Home 3: Sierra, Gaby, and Cathy.* A three-bedroom, single story home occupied by Sierra and Gaby, a female couple, and their roommate, Cathy.

4.2 An Algorithm for Detecting Significant Locations in the Home

Before describing the details of our algorithm, we introduce the following terminology as a matter of convenience:

candidate place: a bounded region proposed as a likely place
merge: the process of combining two candidate places into one candidate place

Candidate places are the primary data element that our algorithm uses. Merging is the primary action performed on these elements.

Algorithm Description The algorithm we use in this work is an agglomerative clustering version of the standard Gaussian Mixture Model (GMM). Our choice was influenced by the fact that we needed to find robust clusters with respect to noisy position readings (which led us toward GMMs) and the fact that we did not want to hand-code the number of places to find (which led us toward an agglomerative process).

To ground our discussion, we first describe the standard GMM algorithm. The GMM algorithm is an iterative process for fitting Gaussian probability distributions to data, consisting of two sub-steps. The pseudo-code is as follows:

1. choose a fixed number of clusters, M

2. initialize these clusters to cover different portions of the data
3. loop until convergence
   (a) for each data point and each cluster, calculate the probability that the cluster generated the data point
   (b) for each cluster, calculate the best Gaussian parameters to explain the data that is strongly associated with it

Usually, the initialization is performed using k-means. We refer to the Gaussian clusters in our algorithm as candidate places. Each candidate place is a three dimensional Gaussian distribution. Two of the dimensions correspond to position and one to time.[3]

Our algorithm augments the basic iterative architecture of the standard GMM with one additional sub-step: an agglomerative step that merges clusters meeting various merging criterion. This additional step occurs after step 3 in the pseudo-code above. We employ three different merging criteria during three phases of our algorithm. At a high level, these phases are (1) merge data points adjacent in space and time into "visits", (2) merge temporally adjacent visits that do not individually meet a minimum duration, (3) merge spatially similar visits into significant locations. Each of these merging steps has an associated parameter that relies on a measure of closeness between candidate places. In describing the algorithm below, we will introduce these parameters, the values we used for the parameters, and the measure of closeness that they rely on.

The algorithm starts by creating one candidate place for each data point. Because each candidate place is initially generated from one data point, the associated Gaussian distribution will have zero variance in the two position dimensions. To combat this issue, we require all candidate places to have a minimum variance along each dimension. Although our position data has sub-meter accuracy, the sensors can give sporadic noise readings that are several meters away. Consequently, we set the minimum variance to be 1.0 for each dimension of the Gaussian distribution.

The choice of minimum variance determines the minimum size of a candidate place. For our choice of parameters, this means that no candidate place can be smaller than about 1 meter in radius. For our data, this is a reasonable constraint. Obviously, the minimum variance values will have to be adjusted for data covering different temporal and spatial scales.

Now, we discuss the three phases of the algorithm and their associated merging steps. In the first phase, we use the following heuristic for merging candidate places: *if two candidate places are temporally adjacent and cover similar positions, they are likely covering data from the same place and should be merged into one visit.* "Similar spaces" is captured in a parameter, λ_1. λ_1 is a threshold on the joint KL-divergence[4] between the two candidate place probability distributions. This heuristic is utilized in most existing place-finding algorithms,

[3] To simplify the calculation performed by our algorithm, we treat each dimension in the Gaussian distribution as independent.

[4] KL-divergence, or Kullback-Leibler divergence, is a calculation of the similarity between two probability distributions. It is written $D_{KL}(P||Q)$, where P and Q are

used to determine when two location readings come from the same place, and is similarly enforced using a threshold on some distance metric between candidate places.

In the second phase, we continue to apply the heuristic used in phase one and also apply the following heuristic: *a place visit should explain a minimum duration of data.* We encode this heuristic in the parameter τ. Again, many place-finding algorithms encode this heuristic in a parameter as well. Generally, it is described as the minimum time that a person must spend in a candidate place before it is treated as a true place.

Finally, in the third phase we apply the following merging heuristic: *if two candidate place visits cover similar spaces but are not temporally adjacent, they are likely encoding two different visits to the same place and should be merged.* This heuristic is captured in the parameter λ_2, which, like λ_1 from phase 1, is a threshold on the joint KL-divergence between candidate places. Many existing place-finding algorithms also utilize this heuristic, and there has even been discussion on different metrics for capturing it numerically [8]. There is an important distinction between this merging step and the previous two. When merging two candidate places, instead of averaging the distributions for each place in the time dimension, we create a new candidate place whose distribution is the union of the two candidate place's distributions. As a result, candidate places take on a slightly new form during phase three: they have a single two-dimensional Gaussian distribution of x and y positions but potentially many Gaussian distributions over time. This allows each candidate place to account for many disjoint visits to the same place.

The space and time complexity of this algorithm depends on both the parameter values and the input data. In the best case, the algorithm will use $O(Nlog(N))$ time and space, where N is the initial number of data points. In the worst case, it will require $O(N^2)$. We implemented the algorithm in Matlab and ran it on a Pentium 4 with 1GB of RAM. The algorithm took between 30 minutes and 2.5 hours to run (we had between 38K and 110K points for each person).

Setting Parameter Values The algorithm has three parameters: λ_1, τ, and λ_2 corresponding to the three different phases of merging. λ_1 encodes the conditions for merging time-adjacent points into candidate places and is constrained by the geometry of the data and the minimum variance values that are chosen for the initial candidate places. We could treat the minimum variance values as additional parameters; however, we found the following steps for adjusting λ_1 provide sufficient control over phase one of the algorithm. The first step is to set the minimum variance values based on the geometry of the data and a common-sense idea of minimum place size. We used 1.0 meter for both position dimensions and 1.0 millisecond for the temporal dimension. The second step is to consider the distribution of joint KL-divergences between the initial candidate

probability distributions over the same space. Since KL-divergence is not symmetric in its arguments, we calculate $D_{KL}(P||Q) + D_{KL}(Q||P)$, the joint KL-divergence.

places. Although not strictly bi-modal, this distribution exhibited a strong peak of values such that choosing $\lambda_1 = 1.0$ merged most of the candidate places that covered similar positions and avoided merging candidates from different places. We fixed $\lambda_1 = 1.0$ for the remainder of this discussion and for all the results presented in section 4.4.

To set the second parameter, τ, which encodes the minimum duration of a visit to a place, we tested many possible values ranging from 30 seconds up to 10 minutes. As discussed in [8], this parameter needs to satisfy conflicting constraints. First, it needs to be short enough that good candidate places are not discarded for not having been visited long enough. Second, it needs to be long enough so that spurious places, created by noisy sensor readings or by brief pauses in an insignificant location, are discarded. Since we are interested in finding places with relatively short visits, like bathrooms, a short minimum duration seemed appropriate. For our data, we found that a τ value of 1 or 2 minutes worked well for most participants.

Finally, λ_2, which encodes the conditions for merging spatially overlapping places during phase three, needs to be set so that candidate places covering approximately the same place are merged while candidate places covering different places are not. The value of this parameter depends on the ratio between the size of single place to the distance between places. Unlike at the city scale, where sensor readings rarely exhibit noise that makes a point intended for one place, e.g. "work", appear near a different place, e.g. "grocery store", our sensor data at the scale of a single home often exhibits noise characteristics that make readings ambiguous as to which place a person is in. Although not a rigorous test, we used two out of the three houses to observe the effects of this parameter (treating the third house as a held-out test set). $\lambda_2 = 2.0$ works well for the data used in this study.

4.3 Measuring the Similarity Between Places

The similarity score that we calculate is derived from the output of our place-finding algorithm. We calculate the score as follows. First, we use each data point's posterior distribution (over candidate places) to label it with its most likely candidate place. In general, the posterior distributions for data points are very peaked, so taking the maximum value is appropriate. Next, using the sequence of data point labels, we calculate the duration of each visit made to each place. For example, if ten successive data points covering 45 seconds all received the same place label, and the preceding and succeeding data points received different labels, then these ten data points constitute a single visit lasting 45 seconds. We also record the time of day that this visit covers. From these visit segments, we calculate a joint distribution over durations (discretized into exponentially growing bins, e.g. 1-2min, 2-4min, 4-8min) and times-of-day (discretized into twenty-four overlapping 5 hour windows) for each place. We claim that these joint distributions roughly characterize the shared attributes of a place, based on the notion that activity patterns follow regularities in terms of their duration and the time-of-day during which they occur. The logical measurement to make

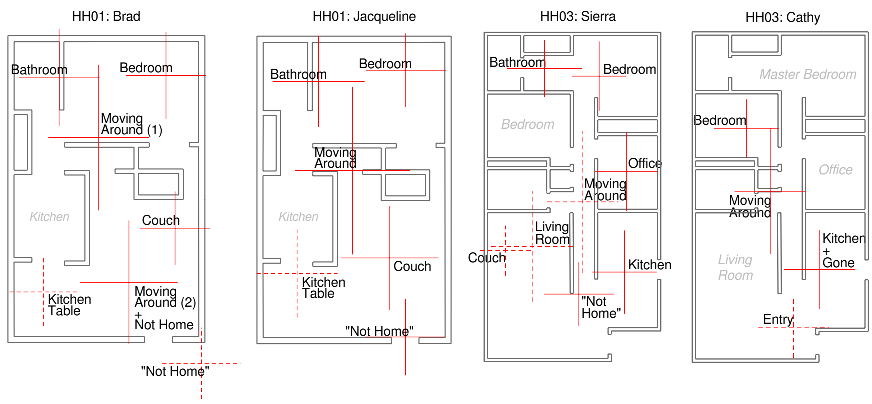

Fig. 1. Places found for several study participants. Each place is centered on the Gaussian distribution found for that place; the length of the cross lines shows one standard deviation in the X and Y directions. Places plotted with dashed lines account for less than 5% of the data, and italic labels show regions not labeled as a place for each person.

between distributions is the joint KL-divergence. Although the KL-divergence between the complete joint probability distributions (over durations and time-of-day) provided reasonable results, we found that using only the distribution over durations resulted in better matchings between the places we found in our data set, and we use this distribution in calculating the similarity scores reported below.

4.4 Results

We ran our place-finding algorithm on the data from each study participant, throwing out places accounting for less than 1% of all data. This resulted in an average of 6.6 places per person (ranging from 4 to 9). In general, each person had one place accounting for "not home" data (participants placed their sensors at a predetermined location, typically by the front door, when they were away from home) and one place accounting for sleeping. The remaining places accounted for activity patterns while participants were at home during the day. Sample places from Households 1 and 3 are shown in figure 1. The fit of these places to the true place in the home, based on experiences in interviews with participants, is discussed in section 4.5.

The similarity measures between all pairs of places were also calculated as described previously. We used these measures to find the optimal many-to-one bipartite matching between the places for each pair of people (matching each place from one person to the "closest" place from each other person). Sample matchings are shown in table 1.

Table 1. Similarity-based place mappings. Italic place-names indicate unexpected mappings.

Within household							
Brad	Table	Living	Gone	*Moving(1)*	Bath	Couch	Bed
Jacqueline	Table	Living	Gone	*Bath*	Bath	Living	Bed
Margaret	Table	*Kitchen*	Couch	Bath	*Gone*	Spare room	
Jack	Table	*Gone*	Couch	Bath	*Kitchen*	Spare room	
Between households							
Sierra	Office	*Bed*	*Living*	Moving	Bath	Couch	*Gone*
Brad	Table	*Moving(2)*	*Gone*	Moving(1)	Bath	Couch	*Bed*

4.5 Evaluation

Evaluation of Place-Finding Algorithm Quantitatively evaluating the "correctness" of places in the home is a particularly difficult process. Systems designed to operate at the city scale are often evaluated based on the percentage of places in the data that they correctly identify. However, this evaluation metric requires that it is possible to identify each important place by hand. For a city, this is straightforward (Hightower et al. [8] had participants carry around a clipboard and write down every place they visited), but places in the home are less well defined, depending on subtle patterns of use. We argue that the "correct" places in the home should not be identified to match a set of *a priori* hand-labeled locations but rather allowed to emerge from the behavior of users, constrained only by the limitations of the characteristics of place that we defined and built into our algorithm. We evaluate the found places by considering how well they match places that would be expected from a taxonomic labeling of people's places (e.g. bedrooms, living rooms, etc.) and by looking for the presence of interesting places that emerge from unexpected patterns of use. We base these evaluations on the qualitative interviews we conducted in each household [2, 28].

Figure 1 illustrates examples of typical and atypical places in Households 1 and 3. For each person, typical places such as couch, kitchen table, bedroom and bathroom often appear. Because these places emerge only when there are significant lengths of time spent in them, the set of labeled places is different for each individual. This is best shown an example from Household 3, where Sierra and Gaby live together and rent their spare room to Cathy. The sets of places for Sierra and Cathy reflect the partitioning of the house between the different people: Sierra's places include the master bedroom, office, living room, and kitchen; and Cathy's places include her bedroom and the kitchen. This distribution of places highlights the difference in space use for the different people, as well as highlighting the kitchen as the primary shared space in the house.

The benefit of allowing other, unexpected, places to emerge, rather than evaluating them with respect to an *a priori* metric can be seen in an example from Household 1. In the main living area of the house, both Brad and Jacqueline reported that there were two significant places: the living room couch and the kitchen table. However, for both Brad and Jacqueline, our algorithm found three places covering portions of the main living area: the living room, the kitchen table, and a third "moving around" area that included both places.[5] Looking at the periods when this larger space was used, we found that it was used during the night of Brad's birthday party. During that time, the two were moving around in the entire main living area of the house *as if it were a single place*. In this case, our place-finding algorithm revealed a significant pattern of place use that would not have been identified in a hand labeling of places and that was not apparent from our interviews. Looking more closely at the sequence of places visited by Jacqueline during the party, we discovered that she began the party in the "moving around place" and as the end approached began to spend more time in the "kitchen table" and "couch" places, suggesting that her use of space became much more localized at the end of the party.

Similar unexpected places are revealed in data from the other households. During the study period, Jack and Margaret in Household 2 had just moved into their apartment and had not yet purchased a bed, sleeping instead on the futon in the main living room of the house each night. As a result, neither labeled the bedroom as a significant place during the study. However, the spare bedroom did show up as a significant place for both Jack and Margaret. In Margaret's case, this was because she used the room to do exercises. In Jack's case, this was because he had an ant farm in the spare bedroom that he would visit periodically throughout the day to take pictures of the ants. (He was creating a time-lapse video of their behavior.) Again, our algorithm revealed patterns of activity and place use that were not immediately apparent upon examination of the architectural space or reported by participants but instead arose from activity.

Evaluation of Similarity Measure The limited amount of in-home location data available makes it difficult to rigorously test our similarity measure. Our goal here is to highlight the potential usefulness of comparing the similarity between places. We evaluate our measure by showing that the resulting mappings between the places of different people are semantically appropriate (e.g., bedrooms are mapped to bedrooms). We also discuss instances where this mapping is unintuitive. Finally, we consider the usefulness of the mappings in capturing common structures by using them to improve predictive models of people's movement patterns in their home.

Table 1 shows the optimal bipartite mappings between the places of different participants. In general, the mappings between the places of two people from the same household are good. For example, a comparison between Brad and Jacqueline fits the expected semantic mapping between places except for

[5] Brad's moving around place was split and partially merged with a "not home" place because his sensor frequently produced readings in the main living room when it was actually near the front door (where he hung his tag while away).

mapping "Moving" to "Bathroom" (the couch to living room mapping is reasonable because Jacqueline does not have a separate couch place). Similarly, the mapping between Margaret and Jack is as expected, except for mapping between "Kitchen" and "Gone". This incorrect mapping can be attributed to sensor noise – Jack and Margaret's kitchen and door (where they hung their tags) were so close together that readings often bounced between them.

Mappings between participants from different households fit with expected mappings in most instances but show more unexpected mappings than within household comparisons. In the mapping between Sierra and Brad, bedrooms are incorrectly matched. However, this mapping also reveals an interesting behavior pattern: Sierra's office and Brad's kitchen table appear to be incorrectly mapped together. But, interviews with both participants revealed that Sierra tended to work in the office and Brad tended to work at the kitchen table more often than eating there. Because our similarity score compares patterns of behavior between places, it captures the fact that, while Sierra's office and Brad's table appear on the surface to be different, the two in fact function as similar places.

The similarity measures between places can also be used to capture and share common patterns of behavior between different people. Many researchers have attempted to predict movement patterns of people between places. For example, Ashbrook and Starner [3] use a second-order Markov model to predict transitions between significant places, Aipperspach et al. [1] use higher-order Markov models to predict transitions, and Liao et al. [12] use a hierarchical activity model to predict transitions. We suggest that good measures of the similarity between places will help to capture "typical" transition patterns between different types of places (e.g., the pattern "people often go from sleeping to the bathroom").

As a method of evaluating the benefit that can be gained by using typical patterns of movement between places to aid in learning movement patterns for an individual in the home, we have built a 3rd-order Markov model over the sequence of place-visits for each individual, as in [1] and [3]. We then build a model that is "bootstrapped" with data from other individuals, using the mappings between places generated based on our similarity measure. This model is built using the Bayesian mixture method implemented in the SRI Language Modeling Toolkit (SRILM) [21].

In evaluating the performance of each model, we use cross-validated perplexity as reported by the SRI toolkit. (A model that perfectly matches the data has a perplexity of 1; models with less than perfect matches have higher perplexity.) Following Mehta et al. [14], we define F_0 to be the area between the learning curve for the first (single person) model and its optimal value (e.g., the area below each dotted line in figure 2). We then define $F_{0|H}$ as the area between the learning curve for the second (multi-person) model and its optimal value (e.g., the area below each solid line in figure 2). Each F measures the performance of its model based on rate of learning and maximum performance, and the *transfer ratio*, $\frac{F_0}{F_{0|H}}$, measures the performance gained when using information about typical movement patterns between different genres of place.

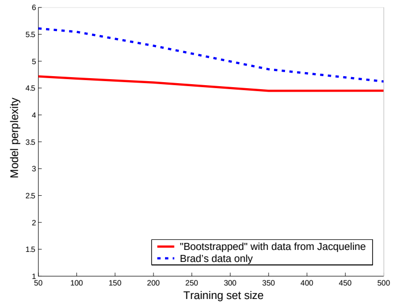

Fig. 2. Learning curves showing transfer between Brad and Jacqueline in Household 1. The dashed line shows the model performance for a model trained only on data from Brad, and the solid line shows the model performance of the same model initialized with data from Jacqueline. The model has a transfer ratio of 1.20.

Figure 2 shows the results of running the predictive models on the data from Brad in Household 1, showing the performance when the model is bootstrapped with data from Jacqueline (in-household transfer). The bootstrapped model has a transfer ratio of 1.197 when compared with the model using only Brad's data, indicating successful transfer of shared movement patterns when using the similarity-based mappings between places.

When evaluating transfer between other study participants, we discovered that our mappings resulted in significantly lower transfer ratios. The primary reason for this is that two people's places do not always have an isomorphic mapping. Hence, requiring a single set of place names to function for two people is a strong assumption to make. In our data set, only Brad and Jacqueline had isomorphic place mappings. We identify two potential solutions to this. First, one could increase the library of known places to include enough cases that an isomorphism is likely to exist between two sets of places. Second, we could allow for probabilistic place name mappings.

5 Design Implications

5.1 Predicting Activities

Several projects are focusing on building models of human activity, which is a difficult process involving the collection and labeling of large histories of data from different types of sensors [6, 19, 22]. However, the shared context of culture can significantly simplify this problem. For example "common sense" repositories (like the Internet) can be automatically mined for information about common tasks [29]. Similarly, since culturally shared notions of place encourage certain activities and not others, an algorithm which can determine how similar a place is to "typical" places of various types could enable context-aware applications to determine the activities most likely to occur in that place. Given a broad enough sample of places from different homes, applications could potentially learn common types of places and then learn typical behaviors for each. Additionally, Ubicomp systems that reference places by name (e.g. [24]) could use names derived from similarity to places with known names to replace hand-labeled names for locations. Currently, the limited availability of accurate location and activity data from multiple homes makes it impossible to fully validate the possibility of learning typical places and associated activities from sensor data. As more data becomes available, we intend to explore more fully the improved activity prediction that can be carried out using shared notions of place.

5.2 Place, Similarity, and Context

An understanding of a user's context is an important aspect of Ubiquitous computing. Dourish [5] describes several problematic assumptions made by traditional views of context. He argues that context is not a stable object, definable separately from the activities of an individual. Rather, he suggests that context is a dynamically defined phenomenon, emerging from and maintained by the patterns of activity collectively carried out by individuals. According to Dourish, our goal as technologists is not "to support particular forms of practice, but to support the evolution of practice."

Our notions of place in the home are in part an attempt to distill meaningful information encoded in people's patterns of movement in a manner compatible with the goals described by Dourish. In particular, both the places that we find and the relationships between them are allowed to emerge entirely from activity. We make few assumptions about the structure of those patterns and instead look for an emergent structure based on people's actions.

The ability to find and compare emergent contextual structures suggests several design possibilities. We consider an example in elder care. One common task in elder care applications is to detect changes in behavior that provide early warning of changes in a person's health [16]. Our similarity metric could be applied to monitor how a single person's place usage changes over time (e.g., by calculating the similarity between the same person's places over two time periods). A low similarity between the sets of places found for an elderly individual may show changes in sleeping patterns or activity levels. These changes in place use could be used to highlight potential health issues. The changes could also be visualized as part of more reflective technologies, such as the ambient health feedback displays developed by Morris [17].

6 Conclusions and Future Work

In this paper, we have presented the first (to our knowledge) algorithm for finding significant places in the home based on location sensor data, and have explored the novel technique of calculating the similarity between places as an aid in understanding the places found by our algorithm. This work is based on the first sub-meter accuracy location traces collected in the home. In order to more fully validate our techniques, particularly in light of the relative uncertainty in the boundaries and definition of place in the home, it will be necessary to continue collecting accurate location traces from a larger set of homes. This will allow us to explore the range of places that exist in the home and to begin to categorize and understand the different types of places that exist across homes, quantifying the work done in [28]. We also plan on collecting and analyzing more diverse types of sensor data. By deploying a broader range of sensors to detect both location and activity, we will be better able to explore the relationship between place and activity and to determine if culturally shared notions of place and appropriate activities can be used to improve activity recognition algorithms.

Currently, the only long-term high-precision location data from the home is in the form of (X,Y) location traces. With current location technology, this data is difficult to collect, requiring a lengthy sensor installation and calibration process. We expect that in-home location traces will become more readily available as location-tracking technologies improve. Another possibility is to explore fingerprint-based place-finding algorithms for the home. Because fingerprint algorithms rely only on recurring patterns of sensor readings, not on detecting the absolute position of each person, the sensor installation process would be greatly simplified. For example, a typical installation may involve only the placement of radio beacons throughout a home instead of the careful installation and calibration of sensors in known positions. Given such fingerprint data for the home, it would be possible to find significant locations and compare their similarity to each other using a variant of our algorithms, at the expense of losing the mapping between the places found and absolute position within the home.

Finally, while we developed our current metric for comparing the similarity between places based on the types of patterns we saw in place usage in the home, we are interested in trying these similarity metrics on place usage at other scales. We plan on applying our similarity measures to existing place data available at the city scale.

References

1. Aipperspach R., Cohen, E., Canny J.: Analysis and Prediction of Sensor Data Collected from Smart Homes. In: Proc. Pervasive '06 (2006)
2. Aipperspach, R., Woodruff, A., Anderson, K., Hooker, B.: Maps of Our Lives: Sensing People and Objects Together in the Home. Tech. Report EECS-2005-22, EECS Department, UC Berkeley (2005)
3. Ashbrook, D., Starner, T.: Using GPS to Learn Significant Locations and Predict Movement Across Multiple Users. In: Personal and Ubiquitous Comp., Vol. 7 (2003)
4. Dey, A., et al.: A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications. In: Human-Computer Interaction, Vol. 16 (2001)
5. Dourish, P.: What We Talk About When We Talk About Context. In: Personal and Ubiquitous Computing, Vol. 8 (2004) 19-30
6. Dragunov, A.N., et al.: TaskTracer: A Desktop Environment to Support Multitasking Knowledge Workers. In: Int'l. Conf. on Intelligent User Interfaces (2005)

7. Harrison, S., Dourish, P.: Re-Place-ing Space: The Roles of Place and Space in Collaborative Systems. In: Proc. CSCW '96 (1996)
8. Hightower, J., et al.: Learning and Recognizing the Places We Go. In: Proc. Ubicomp '05 (2005)
9. Kang, J.H., Welbourne, W., Stewart, B., Borriello, G.: Extracting Places from Traces of Locations. In: Proc. Workshop on Wireless Mobile Applications and Services on WLAN Hotspots, WMASH (2004)
10. Laasonenm K., Raento, M., Toivonen, H.: Adaptive On-Device Location Recognition. In Proc. Pervasive '04 (2004)
11. Leont'ev, A. N.: Activity, Consciousness, and Personality. Prentice-Hall, Englewood Cliffs, NJ (1978)
12. Liao, L., Fox, D., Kautz, H.: Learning and Inferring Transportation Routines. In: Proc. AAAI '04 (2004)
13. Marmasse, N., Schmandt, C.: Location-Aware Information Delivery with Commotion. In: Proc. Symposium on Handheld and Ubiquitous Computing, HUC (2000)
14. Mehta, N., Natarajan, S., Tadepalli, P., Fern, A.: Transfer in Variable-Reward Hierarchical Reinforcement Learning. In: Proc. of the Inductive Transfer workshop at NIPS 2005 (2005)
15. Monteiro, C.: Activity Analysis in Houses of Recife, Brazil. In: Proc. First Int'l. Space Syntax Conf (1997) 20.1-20.13
16. Morris, M., Intille, S.S., Beaudin, J.S.: Embedded Assessment: Overcoming Barriers to Early Detection with Pervasive Computing. In: Proc. Pervasive '05 (2005)
17. Morris, M.: Social Networks as Health Feedback Displays. In: IEEE Internet Computing, Vol. 9(5) (2005) 29-37
18. Oswald, F., et al.: The Role of the Home Environment in Middle and Late Adulthood. In: Wahl, H.-W., et al. (eds.): The Many Faces of Health, Competence and Well-being in Old Age: Integrating Epidemiological, Psychological and Social Perspectives. Springer, Heidelberg (2006)
19. Philipose, M., et al.: Inferring Activities from Interactions with Objects. In: IEEE Pervasive Computing (October 2004) 50-57
20. Rowan, J., Mynatt, E.D.: Digital Family Portraits: Providing Peace of Mind for Extended Family Members. In: Proc. CHI '01 (2001)
21. Stolcke, A.: SRILM - An Extensible Language Modeling Toolkit. In: Proc. of the Int'l. Conf. on Spoken Language Processing (2002)
22. Munguia Tapia, E., Intille, S.S., Larson, K.: Activity Recognition in the Home Setting Using Simple and Ubiquitous Sensors. In: Proc. Pervasive '04 (2004)
23. Trevisani, E., Vitaletti, A.: Cell-Id Location Technique, Limits and Benefits: An Experimental Study. In: Proc. IEEE Workshop on Mobile Computing Systems and Applications, WMCSA (2004)
24. Truong, E., et al.: CAMP: A Magnetic Poetry Interface for End-User Programming of Capture Applications for the Home. In: Proc. Ubicomp '04 (2004)
25. Ubisense, Ltd. Website, http://www.ubisense.net.
26. Wang, J., Canny, J.: End-User Place Annotation on Mobile Devices : A Comparative Study. In: Extended Abstracts, CHI '06 (2006)
27. Wicker, A.: An Introduction to Ecological Psychology. Brooks/Cole Publishing Company, Monterey, California (1979)
28. Woodruff, A., Anderson, A., Mainwaring, S.D., Aipperspach, R.: Portable, But Not Mobile: A Study of Wireless Laptops in the Home. Tech. Report EECS-2006-88, EECS Department, UC Berkeley (2006)
29. Wyatt, D., Philipose, M., Choudhury, T.: Unsupervised Activity Recognition Using Automatically Mined Common Sense. In: Proc. of AAAI '05 (2005)