

User-Dependent Aspect Model for Collaborative Activity Recognition*

Vincent W. Zheng and Qiang Yang

Hong Kong University of Science and Technology, Kowloon, Hong Kong
 {vincentz,qyang}@ijcai.org

Abstract

Activity recognition aims to discover one or more users' actions and goals based on sensor readings. In the real world, a single user's data are often insufficient for training an activity recognition model due to the data sparsity problem. This is especially true when we are interested in obtaining a personalized model. In this paper, we study how to collaboratively use different users' sensor data to train a model that can provide personalized activity recognition for each user. We propose a user-dependent aspect model for this collaborative activity recognition task. Our model introduces user aspect variables to capture the user grouping information, so that a target user can also benefit from her similar users in the same group to train the recognition model. In this way, we can greatly reduce the need for much valuable and expensive labeled data required in training the recognition model for each user. Our model is also capable of incorporating time information and handling new user in activity recognition. We evaluate our model on a real-world WiFi data set obtained from an indoor environment, and show that the proposed model can outperform several state-of-art baseline algorithms.

1 Introduction

With the proliferation of sensor technologies, recognizing human's activities of daily living (ADL) from a series of low-level sensor observations has drawn a lot of research interests in both AI and ubiquitous computing communities. Accurate activity recognition can help us provide various personalized support for many real world applications. For example, [Pollock *et al.*, 2003] used activity recognition to help the elders against the cognitive decline by sending personalized activity reminders. [Eagle and Pentland, 2009] extracted user's eigen-behaviors to help discover community affiliations. Early activity recognition algorithms are based on logic, in which conclusions are deducted from the observations and a number of

"closed world" assumptions [Kautz, 1987]. Bui gave a general model for online plan recognition [Bui, 2003]. Gei and Steedman further showed to use natural language processing to assist plan recognition [Geib and Steedman, 2007]. As extensive sensor data become available, recent activity recognition research starts to focus on using real-time sensor data and learning techniques to recognize the user behaviors. For example, Tapia *et al.* developed a wireless state-change sensor system and used a Naive Bayes classifier to recognize the ADLs [Tapia *et al.*, 2004]. Patterson *et al.* used RFID sensors to capture the abstract object usage and thus use them to recognize user's fine-grained activities with Hidden Markov Models [Patterson *et al.*, 2005]. Liao *et al.* applied a hierarchical Conditional Random Fields to extract a user's locations and activities from GPS data [Liao *et al.*, 2007].

In general, the success of training an activity recognition model relies on having sufficient sensor data from user. However, in the real world, a single user's data are often insufficient due to the data sparsity problem. This is especially true when we are interested in obtaining a personalized activity recognition model. Consider a real-world example when we are trying to build a multi-user activity recognition system with WiFi data. Each user uses her mobile device to record WiFi signal data as she moves in the environment, and annotates the data with some activity label such as "having class". Because the human labeling is expensive and a single user may not foresee all possible WiFi observations in such an open environment, each single user actually does not have enough annotated data to train a personalized activity recognition model for her own. Then we are motivated to ask: as each user's data are insufficient, can we use them together to train a user-dependent model that can provide personalized activity recognition to each user? This problem is not trivial in nature. Most of the previous work does not differentiate the users [Tapia *et al.*, 2004; Lester *et al.*, 2005; Liao *et al.*, 2007]. They treat all the users equally by simply mixing their data in training. However, different users may behave differently given similar sensor observations. For example, a user may visit the coffee shop for meal and the other just enjoys sitting in its outdoor couches to read research paper. These two users probably observe similar WiFi signals, but their activities are quite personalized. This implies that it may not be appropriate to require all the users to share one common, user-independent activity recognizer. In this paper,

*This research is supported by Hong Kong RGC/NSFC N_HKUST624/09, Hong Kong RGC grant 621010 and Microsoft Research Asia MRA10EG01. We also thank the reviewers for their insightful comments.

we propose a user-dependent aspect model to help the users collaboratively build an activity recognition model that can give personalized predictions. Rather than simply pooling multiple users' data together, our model introduces user aspect variables to capture the user grouping information from their data. As a result, for a targeted user, the data from her similar users in the same group can also help with her personalized activity recognition. In this way, we can greatly reduce the need for much valuable and expensive labeled data required in training the personalized recognition model.

Our contributions are summarized as follows:

- We propose a new model to collaboratively utilize different users' data for personalized activity recognition.
- Our model is capable of encoding time information and handling new users in activity recognition.
- We evaluate our model with real-world data and show that it can outperform several state-of-art baselines.

2 Related Work

Our work is different from other multi-user activity recognition work that aims to model concurrent activities among multiple users. For example, Lian and Hsu used a factorial Conditional Random Fields model for joint recognition of multiple concurrent chatting activities [Lian and Hsu, 2009]. Wang et al. proposed a coupled Hidden Markov Model to capture user interactions and recognize multi-user activities at the same time [Wang et al., 2009]. Comparatively, our work does not intend to model such concurrent activities. Besides, our model is able to use the activity data collected at different time by multiple users for personalized activity recognition.

There has been some interesting work that tries to address the data sparsity problem from a different perspective. Rather than aggregating multiple users' data, they focus on utilizing extra knowledge such as human common-sense and unlabeled sensor data. For example, Perkowitz et al. proposed to mine the natural language descriptions of activities from ehov as labeled data, and translated them into probabilistic collections of object terms for training the recognition model [Perkowitz et al., 2004]. In addition to the common-sense knowledge, Wyatt et al. further used the unlabeled RFID sensor data and trained a Hidden Markov Model for activity recognition [Wyatt et al., 2005]. Though in this paper, we do not explore such extra knowledge for our collaborative activity recognition task, we believe it is a promising direction for future research.

Our model is an extension to the standard aspect model [Hofmann and Puzicha, 1999], which is originally proposed to define a generative model for word/document co-occurrences. We use the aspect model to formulate sensor data for activity recognition, and further extend it by introducing user latent aspects to capture the user grouping information. Some closely related concepts to aspect model include probabilistic Latent Semantic Analysis [Hofmann and Puzicha, 1999], Latent Dirichlet Allocation [Huynh et al., 2008], Author Topic Model [Farrahi and Gatica-Perez, 2008]. They are introduced to activity recognition research, and shown to work well to discover implicit activity patterns from

various sensor data. Different from our work, they usually do not explicitly use the activity labels and more focus on data analysis rather than real-time recognition as us.

3 User-dependent Aspect Model

In a WiFi environment, multiple users collect wireless signal data with activity annotations for around a month. The data format is in a set of quads: $\{\langle a_i, u_i, f_i, t_i \rangle | i = 1, \dots, L\}$, where a is an activity, u is a user, and f is a feature observed at time t . In our WiFi case, a feature corresponds to a wireless access point (AP) that our mobile device can detect. A data record quad indicates that a user u is doing an activity a at time t , and meanwhile her wireless device detects some AP f . Our goal is to build a personalized activity recognition model by using these data, so that with a user's WiFi observations at some time, we can predict what she is doing.

3.1 Graphical Model and Its Inference

Our model extends the standard aspect model [Hofmann and Puzicha, 1999; Si and Jin, 2003] by introducing user aspects, as well as time aspects and feature aspects to model personalized activity recognition from time-dependent sensor data. Figure 1 depicts our graphical model. The shadow nodes for user variables u , time variables t , feature variables f and activity variables a are observations. The blank nodes inside the rectangle are latent aspect variables. The user latent aspects $Z_u \in \{z_u^1, z_u^2, \dots, z_u^{D_u}\}$ are discrete variables, indicating D_u user clusters. Our model adopts such a *user-cluster-activity* hierarchy to help the users to collaboratively build an activity recognizer. In contrast to a two-tier *user-activity* hierarchy where each user can only rely on herself to do activity recognition, our model can make the users from a same cluster to contribute together to train the recognizer from their feature and time observations. Therefore, even if some user has limited data to train an activity recognition model, she can still benefit from other similar users in the same group(s). As each user can belong to multiple user clusters at the same time with different probabilities, they actually contribute differently to each user cluster in training the recognition model and consequently get different predictions in real-time recognition. This helps to achieve the model personalization.




Figure 1: User-dependent aspect model.

We also introduce the latent aspects $Z_f \in \{z_f^1, z_f^2, \dots, z_f^{D_f}\}$ and $Z_t \in \{z_t^1, z_t^2, \dots, z_t^{D_t}\}$ to encode the data observations on feature and time. They are used to capture the dependency

between activities and observations, considering that similar feature observations at similar time periods are likely to imply some same activity. Note that these aspects do not necessarily rely on users. We can also take all the users' data as input, and only use them (i.e. Z_f and Z_t) to build a user-independent model for general activity recognition. However, such a user-independent model, as we will show in the experiment, does not perform as well as our user-dependent model. The user latent aspects help to achieve personalization, and that is the reason why we call our model a user-dependent aspect model.

In general, our aspect model is a generative model, which uses the latent aspect variables to explain the observations. It specifies a joint probability of the observed random variables:

$$P(a, u, f, t) = \sum_{Z_u, Z_f, Z_t} P(a, u, f, t, Z_u, Z_f, Z_t), \quad (1)$$

where $P(a, u, f, t, Z_u, Z_f, Z_t)$ is expanded, according to the graphical model, as follows:

$$P(a, u, f, t, Z_u, Z_f, Z_t) = P(Z_u)P(Z_f)P(Z_t)P(u|Z_u)P(f|Z_f)P(t|Z_t)P(a|Z_u, Z_f, Z_t). \quad (2)$$

Here, the user variables u , feature variables f and activity variables a are all discrete in nature, so their conditional probabilities on latent aspects can be modeled easily by multinomial distributions. One exception is the time, which could be continuous. To formulate $P(t|Z_t)$, we discretize the time t with two possible strategies. One is "ByHour", which segments the time into hours. The other is "ByPeriod", which segments it into larger time periods; for example, we can define five periods, including morning (7am~11am), noon (11am~2pm), afternoon (2pm~6pm), evening (6pm~12am) and night (12am~7am). We will compare these two strategies in the experiment to see the impact of time factor.

In inference, for an existing user, we want to predict her activity based on all the access point features detected¹ at some time t . Therefore, our model outputs the activity a^* that has the highest likelihood:

$$a^* = \arg \max_a \prod_f P(a, u, f, t). \quad (3)$$

For a new user, as we do not have any of her data before, we will summarize all the users' predictions to output an activity:

$$a^* = \arg \max_a \sum_u \prod_f P(a, u, f, t). \quad (4)$$

3.2 Model Training

For model training, we can use the Expectation Maximization (EM) algorithm to get the maximum likelihood estimation². At E-step, for each data example $\langle a_i, u_i, f_i, t_i \rangle$, we compute

$$P(Z_u, Z_f, Z_t | a_i, u_i, f_i, t_i) = \frac{P(a_i, u_i, f_i, t_i, Z_u, Z_f, Z_t)}{\sum_{Z_u, Z_f, Z_t} P(a_i, u_i, f_i, t_i, Z_u, Z_f, Z_t)}. \quad (5)$$

At M-step, given L data examples, we compute

$$P(Z_u) = \frac{1}{L} \sum_i \sum_{Z_f, Z_t} P(Z_u, Z_f, Z_t | a_i, u_i, f_i, t_i), \quad (6)$$

¹We follow [Yin *et al.*, 2004] to assume AP independence.

²Derivation details are skipped here for space, interested readers can refer to [Hofmann and Puzicha, 1999; Si and Jin, 2003].

$$P(u|Z_u) = \frac{\sum_{i:u_i=u} \sum_{Z_f, Z_t} P(Z_u, Z_f, Z_t | a_i, u_i, f_i, t_i)}{L \times P(Z_u)}, \quad (7)$$

$$P(a|Z_u, Z_f, Z_t) = \frac{\sum_{i:a_i=a} P(Z_u, Z_f, Z_t | a_i, u_i, f_i, t_i)}{\sum_i P(Z_u, Z_f, Z_t | a_i, u_i, f_i, t_i)}. \quad (8)$$

Analogous to Eq.(6), we can get $P(Z_f)$ and $P(Z_t)$; and similarly, analogous to Eq.(7), we can get $P(f|Z_f)$ and $P(t|Z_t)$. Due to space limit, we skip the details here.

For model complexity, let's denote the number of activity as N_a , the number of users as N_u , the number of features as N_f . Together with D_u user aspects, D_f feature aspects and D_t time aspects, we totally need to maintain $N = (N_a \times D_u \times D_f \times D_t + N_u \times D_u + N_f \times D_f + N_t \times D_t + D_u + D_f + D_t)$ variables for our model. The time complexity for each EM iteration is as follows: at E-step, we need to update the probability of $P(Z_u, Z_f, Z_t | a_i, u_i, f_i, t_i)$ for each data example by summing over the latent aspects, therefore the cost is $O(L \times D_u \times D_t \times D_f)$; at M-step, we can amortize the summation over L samples as $O(L \times D_u \times D_t \times D_f)$ and after one data scan update the N variables by $O(N)$.

4 Experiments

In our experiments, we have 13 users collecting the WiFi data for around a month, basically in a university area. Their mobile devices sniffed the WiFi signals roughly every 10 minutes when it was power on. The users annotated their WiFi data with eight possible activities from time to time, and we further preprocess the data by data segmentation and label parsing. On average, each user has around 1,150 data examples, and in total we observed 2,912 different access points (i.e. features) in this dataset. Some activity and user statistics is shown in Figure 2. For each user, we split the first half of her data in chronological order as training data, and the other half as test data. We measure the average accuracy among all the users at each trial, and report the average accuracy of three trials through the experiments. Without special notification, our model uses the "ByPeriod" time segmentations, and the model parameters are set as: $D_u = 3$, $D_t = 2$ and $D_f = 20$.

4.1 Activity Recognition for Existing Users

In this experiment, we gather the same percent of training data from each user and fit them into our model for training. Then, we test each user and give the average accuracy. We also vary the number of training data and see the performance change. We employ three baselines for comparison: (1) "Single" baseline, which uses each single user's data for training and maintains a personalized activity recognition model for each user; (2) "Merged" baseline, which pools all the users' data together and trains a general activity recognition model for all the users. For these two baselines, we use the sophisticated Conditional Random Fields [Vail *et al.*, 2007; Liao *et al.*, 2007] as the classifier model. These two baselines also take the time information as a feature input for fair comparison; (3) "MostFreq" baseline, which uses the most frequent activity based on training data as the prediction.

As shown in Figure 3(a), as the number of training data increases, our model is consistently better than the baselines.




Figure 2: Data statistics.




Figure 3: System performance.

The “MostFreq” baseline’s performance keeps unchanged, as the found most frequent activity is always “doing research” as training data size increases. The baseline’s comparatively poor performance shows that a simple solution may not be adequate for this complex task. For the other two baselines, we notice that, when the training data size is small, our model’s improvement over the baselines seems small as well; but as the training data size increases, the improvement increases quickly. This is because with limited training data, the learned user clusters may not be as accurate as that with enough training data. Another interesting observation is that, when there is more training data, the “Single” baseline seems to outperform the “Merged” baseline. It implies that each user’s unique activity pattern can be better preserved by personalized activity recognition given enough training data.

Percent of training data	User-dependent (our model)	User-independent (variant of our model)
20%	0.56 ± 0.02	0.52 ± 0.01
40%	0.57 ± 0.02	0.53 ± 0.02
60%	0.62 ± 0.03	0.58 ± 0.04
80%	0.65 ± 0.01	0.60 ± 0.02
100%	0.71 ± 0.00	0.64 ± 0.03

Table 1: Impact of the user latent factors [acc $\pm$ std].

We also study the impact of user latent aspects to our model, so that we can understand how much benefit this user-dependent solution brings to personalized activity recognition. We employ a variant of our model as the baseline, named user-independent aspect model for comparison. This baseline takes all the users’ data as input. It is a simplified

version of our model with the user latent factor Z_u and user variable u removed. In other words, this baseline only consists of feature and time latent aspects, without encoding any user information, and thus it is user-independent. As shown in Table 1, our user-dependent model consistently outperforms the baseline. It proves that, the user aspects are crucial to the personalized activity recognition task. In addition, we also tested the case when this user-independent baseline only takes a single user’s data as input for training to mimic personalized activity recognition. On average, we observed a 15% performance lift of our model over such a baseline. This shows that our model can well incorporate multiple users’ data.

4.2 Activity Recognition for New Users

In this experiment, we use a leave-one-user-out strategy, with which we hold out one user only for testing and the other users for training. Then we rotate on each user and report an average accuracy in Figure 3(b). As there is no training data for the test user, the “Single” baseline does not work anymore. We compare our model with the “Merged” baseline which requires all the users to share a same activity recognizer. As can be seen in the figure, our model’s performance is comparable with the baseline. This is because when there is no training data for a new user, our model has no way to capture her grouping information. By summing up the prediction probabilities over the users (c.f. Eq.(4)), our model thus gives a comparable prediction result with the “Merged” baseline of pooling all the data together.

4.3 Impact of Time and the Model Parameters

We study the impact of two time segmentation strategies: “ByHour” and “ByPeriod”. As shown in Figure 3(c), two

strategies have comparable performances, though when the training data size is small, the “ByPeriod” strategy seems better. This may be because “ByPeriod” strategy summarizes the reasonable time segments with fewer parameters, which are favored given limited training data. We also observe that, the model performance seems not very sensitive to the latent time aspect cardinality D_t .

We study the impact of the model parameters on user and feature group numbers. In this experiment, we use 60% of the training data. We vary the model parameter pairs (D_u, D_f) , and report the averaged accuracy in the recognition task with existing users. As shown in Table 2, when the user cluster numbers are reasonably small, e.g. $D_u = 3$ or 5 compared with the total user number as $N_u = 13$, the model performance is stable. But when the user cluster number is too big, e.g. $D_u = 10$, the performance could possibly drop due to inappropriate clustering. For the number of feature clusters D_f , as the number of observed APs in our data is big, we see the performances for $D_f = 20$ and $D_f = 50$ are comparable. We also observe similar patterns in recognizing new users.

<i>Parameters</i>	<i>Accuracy</i>
$D_u = 3, D_f = 20$	0.62 ± 0.02
$D_u = 5, D_f = 20$	0.60 ± 0.01
$D_u = 5, D_f = 50$	0.59 ± 0.02
$D_u = 10, D_f = 20$	0.57 ± 0.03

Table 2: Impact of the model parameters on D_u and D_f .

4.4 Discussion

When our model works and when it may not work? These are the questions that we are curious about. We plot our model’s performance on each user in Figure 4. We used 60% of training data, and the baselines are described as above.

In the figure, we can observe our model improvements on most users over the baselines. Our model especially works well on user 3, user 5 and user 11 compared with the baselines. Let us take user 11 as an example for analysis. The “Single” baseline does not work well in this user’s case, because in the test data, we have many access point features unseen in the training data. Therefore, the “Single” baseline, since it only uses a single user’s data for training, cannot handle well the unseen observations and give many incorrect predictions. Comparatively, our model can work better because it can benefit from other users’ data which may contain some unseen access points. Interestingly, the “Merged” baseline also does not work well in this case, though it uses other users’ data. The test data show that user 11 often went to the campus cafe area (which has cozy couches and coffee) in the evening time to do research. But the “Merged” baseline predicts the activity in the cafe area as “having meal”, as most of the users seen in the dataset annotated “having meal” there at that time. While our model correctly gives the prediction of “doing research”, because it figures out that one of his similar user (i.e. user 4) sometimes also did the same thing in the cafe. This pattern is preserved to correct the predictions as “doing research” rather than “having meal”. This case shows that, our model works when: 1) there are multiple users’ data

available; and especially 2) when similar users share the similar activities given similar observations.

There are also some cases when our model may not work well, such as on user 7 and user 8. Though the reasons why our model does not work well on both users appear slightly different, there is one thing in common. That is, our model’s underlying assumption which assumes a group of similar users would do similar activities given some similar observations is violated in these two cases. For example, some of user 7’s test data indicate his “having leisure” in some office. Our model predicts the activities as “doing research” since one of his similar user usually worked in that office area. This activity inconsistency among the users leads to the performance drop. Beside, user 8 is shown to be often “having leisure” in her residence area. But in user 8’s training data, the most frequent activities that happens in her residence area are “doing research” and “others”. Her similar users seem to have consistent activity patterns in training and test data. Consequently, this behavior difference in training/test data also makes the model perform poorly. Notice that in both user 7 and user 8, the “Merged” baseline seems better, as it suffers less from such behavior differences by considering all the users equally.

5 Conclusion

In summary, we study how to collaboratively use different users’ sensor data for personalized activity recognition. We introduce a user-dependent aspect model which formulates the user grouping information with latent user aspects. Therefore, each single user can benefit from other similar users’ data. Our model can also use the time information and provide activity recognition to a new user. We test our model with our real-world WiFi-based activity recognition system. For existing users, we show average performance lift (w.r.t. different training data sizes) of 11% over the baseline with pooling all the users’ data together and 13% over the baseline with each single user’s data. For activity recognition with new users, our model also gives comparable performances with the baseline. We also give case studies to help understand the conditions for our model to work.

For future work, we are interested in extending our model to online update with new observations. We also want to explore using extra knowledge such as common-sense knowledge and unlabeled sensor data to help improve the system.

References

- [Bui, 2003] Hung Hai Bui. A general model for online probabilistic plan recognition. In *Proc. of the 18th International Joint Conference on Artificial Intelligence (IJCAI ’03)*, pages 1309–1318, 2003.
- [Eagle and Pentland, 2009] Nathan Eagle and Alex Pentland. Eigenbehaviors: Identifying structure in routine. In *Behavioral Ecology and Sociobiology*, pages 1057–1066, May 2009.
- [Farrahi and Gatica-Perez, 2008] Katayoun Farrahi and Daniel Gatica-Perez. What did you do today?: discovering daily routines from large-scale mobile data. In *Proc.*




Figure 4: Performance on each user.

of the 16th ACM International Conference on Multimedia (MM '08), pages 849–852, 2008.

[Geib and Steedman, 2007] Christopher W. Geib and Mark Steedman. On natural language processing and plan recognition. In *Proc. of the 20th International Joint Conference on Artificial Intelligence (IJCAI '07)*, 2007.

[Hofmann and Puzicha, 1999] Thomas Hofmann and Jan Puzicha. Latent class models for collaborative filtering. In *Proc. of the 16th International Joint Conference on Artificial Intelligence (IJCAI '99)*, pages 688–693, 1999.

[Huynh et al., 2008] Tâm Huynh, Mario Fritz, and Bernt Schiele. Discovery of activity patterns using topic models. In *Proc. of the 10th International Conference on Ubiquitous Computing (UbiComp '08)*, pages 10–19, 2008.

[Kautz, 1987] Henry Kautz. *A Formal Theory of Plan Recognition*. PhD thesis, University of Rochester, 1987.

[Lester et al., 2005] Jonathan Lester, Tanzeem Choudhury, Nicky Kern, Gaetano Borriello, and Blake Hannaford. A hybrid discriminative/generative approach for modeling human activities. In *Proc. of the 19th International Joint Conference on Artificial Intelligence (IJCAI '05)*, pages 766–772, 2005.

[Lian and Hsu, 2009] Chia-Chun Lian and Jane Yung-jen Hsu. Probabilistic models for concurrent chatting activity recognition. In *Proc. of the 21st International Joint Conference on Artificial Intelligence (IJCAI '09)*, 2009.

[Liao et al., 2007] Lin Liao, Dieter Fox, and Henry A. Kautz. Extracting places and activities from gps traces using hierarchical conditional random fields. *International Journal of Robotics Research (IJRR)*, 26(1):119–134, 2007.

[Patterson et al., 2005] Donald J. Patterson, Dieter Fox, Henry A. Kautz, and Matthai Philipose. Fine-grained activity recognition by aggregating abstract object usage. In *Proc. of the 9th IEEE International Symposium on Wearable Computers (ISWC '05)*, pages 44–51, 2005.

[Perkowitz et al., 2004] Mike Perkowitz, Matthai Philipose, Kenneth Fishkin, and Donald J. Patterson. Mining models

of human activities from the web. In *Proc. of the 13th international conference on World Wide Web (WWW '04)*, pages 573–582, 2004.

[Pollack et al., 2003] Martha E. Pollack, Laura E. Brown, Dirk Colbry, Colleen E. McCarthy, Cheryl Orosz, Bart Peintner, Sailesh Ramakrishnan, and Ioannis Tsamardinos. Autominder: an intelligent cognitive orthotic system for people with memory impairment. *Robotics and Autonomous Systems (RAS)*, 44(3-4):273–282, 2003.

[Si and Jin, 2003] Luo Si and Rong Jin. Flexible mixture model for collaborative filtering. In *Proc. of the 20th International Conference on Machine Learning (ICML '03)*, pages 704–711, 2003.

[Tapia et al., 2004] Emmanuel M. Tapia, Stephen S. Intille, and Kent Larson. Activity recognition in the home using simple and ubiquitous sensors. In *Pervasive Computing*, pages 158–175, 2004.

[Vail et al., 2007] Douglas L. Vail, Manuela M. Veloso, and John D. Lafferty. Conditional random fields for activity recognition. In *Proc. of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS '07)*, pages 1–8, 2007.

[Wang et al., 2009] Liang Wang, Tao Gu, Xianping Tao, and Jian Lu. Sensor-based human activity recognition in a multi-user scenario. In *Ambient Intelligence*, volume 5859, pages 78–87, 2009.

[Wyatt et al., 2005] Danny Wyatt, Matthai Philipose, and Tanzeem Choudhury. Unsupervised activity recognition using automatically mined common sense. In *Proc. of the 20th National Conference on Artificial Intelligence (AAAI '05)*, pages 21–27, 2005.

[Yin et al., 2004] Jie Yin, Xiaoyong Chai, and Qiang Yang. High-level goal recognition in a wireless lan. In *Proc. of the 19th national conference on Artificial intelligence (AAAI '04)*, pages 578–583, 2004.