

Activity Recognition in Multi-User Environments Using Techniques of Multi-label Classification

Alaa Alhamoud
Multimedia Communications
Lab
TU Darmstadt
Darmstadt, Germany
Alaa.Alhamoud@kom.tu-
darmstadt.de

Vaidehi Muradi
Multimedia Communications
Lab
TU Darmstadt
Darmstadt, Germany
Vaidehi.Muradi@kom.tu-
darmstadt.de

Doreen Böhnstedt
Multimedia Communications
Lab
TU Darmstadt
Darmstadt, Germany
Doreen.Boehnstedt@KOM.tu-
darmstadt.de

Ralf Steinmetz
Multimedia Communications
Lab
TU Darmstadt
Darmstadt, Germany
Ralf.Steinmetz@kom.tu-
darmstadt.de

ABSTRACT

Activity recognition represents the cornerstone in realizing intelligent services such as energy conservation and ambient assisted living in smart environments. The problem statement of most activity recognition research assumes that only mutually exclusive activities occur in smart environments. The majority of research projects in this field focus on single-user environments where only one user performs a single activity at a given time. Such solutions are not applicable in real-world scenarios where multiple users reside in a home performing co-temporal activities. Our work addresses the problem of activity recognition in multi-user environments by utilizing the techniques of multi-label classification. It is based on a multi-label activity recognition dataset which we collected by deploying appliance-level power sensors as well as environmental sensors in a two-person apartment. In this dataset, a feature vector of sensor readings can have more than one label indicating the occurrence of more than one activity at a given time. In this work, we show that recognizing activities in smart environments can be achieved solely based on fine-granular power consumption data and without the need for installing any other sensing modality. Moreover, we prove that extracting and utilizing dependency relations between concurrent activities as well as temporal relations between

subsequent activities provide a crucial enhancement of the predictive accuracy of activity recognition models.

ACM Classification Keywords

I.5.2 PATTERN RECOGNITION: Design Methodology–Classifier design and evaluation–Pattern analysis; I.5.4. PATTERN RECOGNITION: Applications

Author Keywords

Activity Recognition; Smart Environments; Multi-Label Classification.

INTRODUCTION

Activity recognition in smart environments provides ways to improve several aspects of users' life. It is applied in many fields such as elderly care [13], health care, energy conservation [2] [10] [3] and so on. Major research projects include body wearable sensors that help to provide elderly care and health care [8]. Wearable body sensors along with environmental sensors give more precise information in localization and recognition of a person's activity. However, it is always a case of discomfort to carry gadgets stuck to the body all the time. The scope of this paper confines to using only power sensors in order to avoid this discomfort to the users.

In recent times, wide research has been done on activity recognition in single-user environments. Nonetheless, the common scenario in real world is that multiple users live in a common place. This paper aims at building a model for multi-user activity recognition which is able to detect the current activities of more than one user. Considering the multi-user scenario, there is scope for a user to perform more than one activity at a given time. For example, a user could be working at PC, later go off to cooking while listening to music and come back to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IoT'16, November 07-09, 2016, Stuttgart, Germany

© 2016 ACM. ISBN 978-1-4503-4814-0/16/11...\$15.00

DOI: <http://dx.doi.org/10.1145/2991561.2991563>

working at PC. The second user could be watching movie and then decide to make tea and come back to watching movie.

In our work, we propose a new approach for building a multi-user activity recognition model based on the techniques of multi-label classification [12]. We identify dependency relations between concurrent activities as well as temporal relations between subsequent activities. Based on a comprehensive evaluation study, we show the essential role of these relations in enhancing the predictive performance of multivariate activity recognition models.

The rest of this paper is organized as follows: We present a group of research projects in the field of activity recognition in multi-user environments in Section 2. We introduce the dataset used in this work along with the process of feature extraction in Section 3. Our first approach for building a model for activity recognition in multi-user environments is presented in Section 4. We introduce our approach for identifying dependency relations between concurrent activities and for building a multivariate activity recognition model in Section 5. We conclude the paper in Section 6.

RELATED WORK

Several research projects are available in the field of activity recognition although with limited assumptions [4][5][7][14][11]. A single user performing only one activity at a given time is a common assumption among many research projects. Few projects have worked on recognizing multiple activities occurring parallel in a multi-user environment.

The first work we discuss is conducted by Crandall et al. [5]. It primarily focuses on addressing the two main problems of tracking and identifying multiple residents in a home. To know which activity belongs to which user, differentiating between the behaviors of users is vital. The temporal features such as hour of the day, part of the day, day of the week, and part of the week (weekend or weekday) are used to study users' behaviors. Motion sensors and user interaction with lighting devices are used to record events. The data gathered consisted of date, time, sensor serial number, event message and annotated class which resembles the resident ID. Bayesian classifier was used to model the data. Evaluation of models that include sensors with temporal features shows good results in differentiating the users. The dataset was randomly split into training and testing sets. The classifier was trained on 90% and tested on 10% of each class. Although this scenario even though considers multiple users, it does not consider parallel activities.

Doryab et al. [7] proposed an approach that considers joint recognition of activities in clinical work. Multiple wearable and embedded sensors are placed in the operating rooms. Raw data is collected from these sensors. Base activities are recognized from this data. Apriori algorithm [1] is used to mine frequent patterns from the dataset. These frequent patterns give insight into which base activities occur together. Thus, artificial joint activities can be formed. These joint activities are used to transform multi-label data into single-label data. A virtual evidence boosting algorithm is used to capture the temporal dependencies between actions. The data used in this work was recorded from 10 laparoscopic surgeries. It is

estimated that 6 clinicians participate in a surgery. Sensors were placed to know the location of clinicians, the location of objects and their use by different clinicians. All the sensor values were transformed to binary values. It was observed that more than 70% of the data had more than 2 activities. A Conditional Random Field (CRF) was used to model the data. Initially, model *A* was built that included all the combinations of joint activities. This model was very slow and showed poor performance. This is due to the huge number of joint activities that were considered as single labels. In the next step, the anaesthesia and operating team were separated. Model *B* was trained on each team. This model showed better performance than the previous one. However, the concurrent activities occurring between the two teams could be missed out. To overcome this problem, an additional model *C* was built and trained on these concurrent activities occurring between the two teams.

Wu et al. [14] proposed a method for joint recognition of activities in a multi-user home. The authors used *House_n* dataset which was provided by the Massachusetts Institute of Technology (MIT). The data was recorded everyday from 9 AM to 1 PM. The sensors used in this project were switch, light, and current sensors. The dataset was later annotated with labels based on ground truth of multiple activities and location information provided by the users and sensors respectively. The availability of data was less because the data was collected for only four hours everyday. Therefore, the researchers focused mainly on location information for recognizing activities. The dataset was divided into 18 parts each of 10 minutes. 89 activities were clustered into 6 classes. These 6 classes could happen to overlap with each other thus enabling multi-user activity recognition. A Factorial Conditional Random Field (FCRF) was used to build the data model.

Compared to previous projects, this work aims at building a multivariate activity recognition model that is able to recognize several activities of multiple users simultaneously. This is achieved without the need for converting the data into single-label as done by other projects that combined concurrent activities into joint activities.

DATASET, DATA PREPROCESSING AND FEATURE EXTRACTION

We collected the multi-user dataset used in this work from an apartment with two students residing in it. The collection spanned for a duration of 23 days. Pikkerton smart energy meters ¹ were attached to power devices. The devices that were monitored are listed in Table 1. These sensors collected measurements in terms of current, voltage, frequency and power for each device. Fibaro environmental multi-sensors ² were placed in two locations, namely sitting room and corridor. These locations are depicted in the layout as shown in Figure 1. Each Fibaro sensor measures temperature in Celsius, light intensity in LUX and motion as Boolean (motion or no motion). Both types of sensor nodes provided new measurements every 28 seconds. A smartphone enabled both users to report

¹<http://www.pikkerton.com/>

²<http://www.fibaro.com/>

their current activities. These activities are listed in Table 2. We monitor a set of 11 activities for each user. When the user performs an activity which does not belong to the list of monitored activities, he has to provide *Ignore* as a feedback. The collected dataset includes 335000 sensor readings and 677 user's feedback.

Lamp	Monitor
PC-User1	PC-User2
Oven	Stove
Vacuum Cleaner	Sound System
TV	Water Heater

Table 1. List of monitored appliances

Eating	Reading
WorkingAtPc	WatchingMovie
WatchingTV	Cooking
Cleaning	ListenToMusic
Sleeping	OutOfHome
MakingTea	Ignore

Table 2. List of monitored activities




Figure 1. Placement of Fibaro sensors in the apartment

Feature Extraction

Due to the high sampling rate, it is recommended to reduce the dataset size and extract the important discriminative features that help distinguishing the different activities. Therefore, a windowing technique as specified by Alhamoud et al. [2] was followed. The feature extraction process consists of the following steps:

1. The start time is considered for current activity.
2. The start time of next activity is taken as an end time for the current activity.
3. The time difference between start and end time is taken.
4. It is further divided into 2-minute slots. A time slot length of 2 minutes has been chosen based on the observations obtained from our previous work described in [2].

5. In each time slot, we take the maximum reading of each sensor as a representative feature for this slot. This is clarified in Eq.(1) which shows how to construct feature vectors.
6. The timestamp represented by the hour at which an activity is being performed is considered as a part of the feature vector as well.

$$F(t) = \langle S_1(t), S_2(t), S_3(t) \dots S_n(t), H \rangle \quad (1)$$

Where:

- n : number of sensors
- $S_i(t)$: the maximum value for sensor i in timeslot t
- H : Hour

7. A single user can perform more than one activity at a given time. Hence, a user can be tagged with more than one label. In order to apply multi-label classification on this dataset, the activities are converted into binary labels as shown in Eq.(2). These binary labels indicate if the activity occurs or not for given sensor values.

$$L(t) = \langle L_0, L_1, L_2 \dots L_q \rangle \quad (2)$$

Where:

- q : number of labels
- $L(t) \in \{0, 1\}^q$

8. The dataset consisting of feature vectors along with activities as labels can be represented as follows:

$$Dataset = \begin{bmatrix} F(1), \langle L_1 0 \dots L_1 q \rangle \\ F(2), \langle L_2 0 \dots L_2 q \rangle \\ F(3), \langle L_3 0 \dots L_3 q \rangle \\ \dots \dots \dots \\ F(m), \langle L_m 0 \dots L_m q \rangle \end{bmatrix}$$

As mentioned before, we monitor a set of 11 activities for each user which results in 11 labels. The process of feature extraction is applied on the activities of both users. Thus, two datasets are obtained, namely *User 1* and *User 2* datasets. These two datasets are then merged based on the timestamp to form a unified dataset for both users. Moreover, examples of activities *WatchingTV* and *ListenToMusic* are removed for *User 1* and examples of activities *Reading* are removed for *User 2* due to very few number of examples. In order to obtain training and testing sets, we divided the dataset based on the dates. The days that are placed in training set were not repeated in testing set. Non-overlapping days were used to train and test such that the model can be open to real-time scenarios.

BINARY RELEVANCE AS A MULTI-LABEL CLASSIFICATION TECHNIQUE

In this section, we present our first approach for building an activity recognition model for multi-user environments. It is based on the binary relevance (BR) [12] as a problem

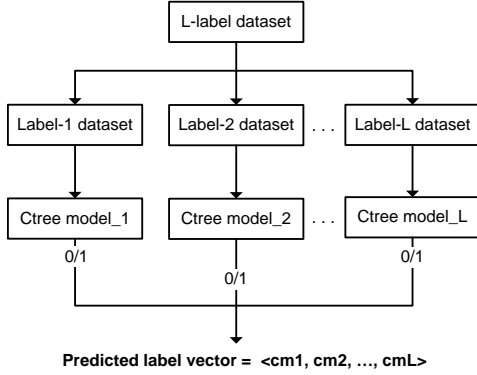


Figure 2. Multi-label data transformation into single-label data

transformation method used in the context of multi-label classification. BR transforms a multi-label classification problem into a set of single-label classification problems by building a separate dataset for each individual label. Any classical single-label classifier can be used to solve the resulting single-label classification problems. We use the algorithm of conditional inference tree (ctree) [9] as a base classifier for our binary relevance approach.

We build a ctree model for each single-label dataset. The response of this tree is the label to be predicted while its predictor variables are the sensor readings comprising the feature vector as shown in Eq.(1). Time as a feature plays an important role in differentiating between activities. Therefore, we incorporate the timestamp of an activity represented by the hour at which this activity is happening into the feature vector. The workflow of BR technique is shown in Figure 2. We conduct an exhaustive evaluation study in order to evaluate the predictive performance of binary relevance models. We start our evaluation in Section 4.1 by considering all deployed sensors, namely power and environmental sensors as parts of the feature vector. In Section 4.2, we study the effect of temporal relations between subsequent activities on the predictive performance of binary relevance models. As our main goal is to recognize users’ activities solely based on fine-granular measurements of power consumption, we analyze the effect of excluding environmental sensors from feature vectors on the overall predictive performance of binary relevance models in Section 4.3.

Power and Environmental Sensors

This section discusses in detail the predictive performance of binary relevance models. In this evaluation setting, we consider the feature vectors to contain values of power and environmental sensors combined with the hour at which an activity has occurred as represented by Eq. (1). Figure 3 and 4 show the f-measure values achieved by binary relevance models in recognizing activities of *User 1* and *User 2* respectively. The activities followed by “.x” are for *User 1* and by “.y” are for *User 2*. As we can see from both figures, binary relevance models were able to predict the activities of both users with good predictive performance. However, they were




Figure 3. F-measure values of environmental-power binary relevance models with regard to *User 1*




Figure 4. F-measure values of environmental-power binary relevance models with regard to *User 2*

not able to predict any instance of the activity *MakingTea*. This is due to the very few number of instances for this activity in comparison with other activities. In Section 5, we introduce a solution for this problem by studying and identifying dependency relations between the different activities.

Temporal Relations

Users tend to follow a certain routine in performing their daily activities. In such a routine, they are used to do their activities in a sequence that repeats itself everyday. In this section, we study the effect of such temporal patterns on the predictive performance of activity recognition models. To achieve this goal, we transform the dataset so that the previous activity for a given time slot is added as a feature. The dataset will then consist of instances such that each instance consists of a feature vector comprised of values of power and environmental sensors, timestamp and activity previously performed by the user as shown in Eq. (3).

$$F(t) = \langle S_1(t), S_2(t), S_3(t) \dots S_n(t), H, prevAc \rangle \quad (3)$$

Where:

- n : number of sensors
- $S_i(t)$: the maximum value for sensor i in time slot t
- H : hour




Figure 5. The improvement in f-measure after encoding temporal relations as extra features for the activities of *User 1*




Figure 7. Comparison between temporal and non-temporal models in terms of average f-measure, accuracy, precision and recall




Figure 6. The improvement in f-measure after encoding temporal relations as extra features for the activities of *User 2*

- *prevAc*: previous activity

Figures 5 and 6 show the improvement achieved in f-measure values for the activities of both users after the inclusion of temporal patterns into the feature vectors. As we can see from both figures, activities such as *Eating*, *Cooking*, *OutOfHome* and *ListenToMusic* have seen a very good improvement for both users. However, it is still not possible for the model to predict any instance of the activity *MakingTea*. Figure 7 shows an improvement of 24%, 19%, 13% and 3% achieved in the average values of precision, f-measure, recall and accuracy respectively. Average values have been calculated with respect to all activities of both users.

Power Sensors

Our main goal in this work is to build a model that recognizes users' activities in multi-user environments solely based on fine-granular sensing of power consumption and without the need for deploying any other sensing modality. Therefore, this section studies the effect of excluding the values of environmental sensors from feature vectors on the overall performance of binary relevance models. Figures 8 and 9 compare the f-measure values of power-environmental and power-only models. Both figures show that the exclusion of environmental sensors has not caused any significant decrease in f-measure values for both users. It has even caused an in-




Figure 8. Comparison between power-only and power-environmental models in terms of f-measure values for activities of *User 1*

crease in f-measure values for the activities *OutOfHome* and *Eating* of *User 2*. However, the model is still not able to recognize the activity of *MakingTea* due to its very few number of instances.

In a multi-user environment, it is common that concurrent activities may occur between the users or that a single user may perform multiple activities at a given time. This leads to the existence of a dependency pattern in which certain activities occur or never occur together most of the time. A user may eat and watch TV at the same time. In case of two users, they usually tend to eat, watch tv and cook at the same time. Binary relevance algorithm has the main disadvantage of building a separate model for each individual label ignoring any useful dependency information that may exist between activities. To improve the predictive performance of our models, it is important to study and identify such dependency patterns. Therefore, we introduce in Section 5 our approach for identifying and extracting dependency patterns between activities using the algorithm of conditional inference trees.

LABEL DEPENDENCY

In this work, we identify and extract two different types of unconditional label dependency [6]. Two activities are unconditionally dependent on each other i.e. correlated if they

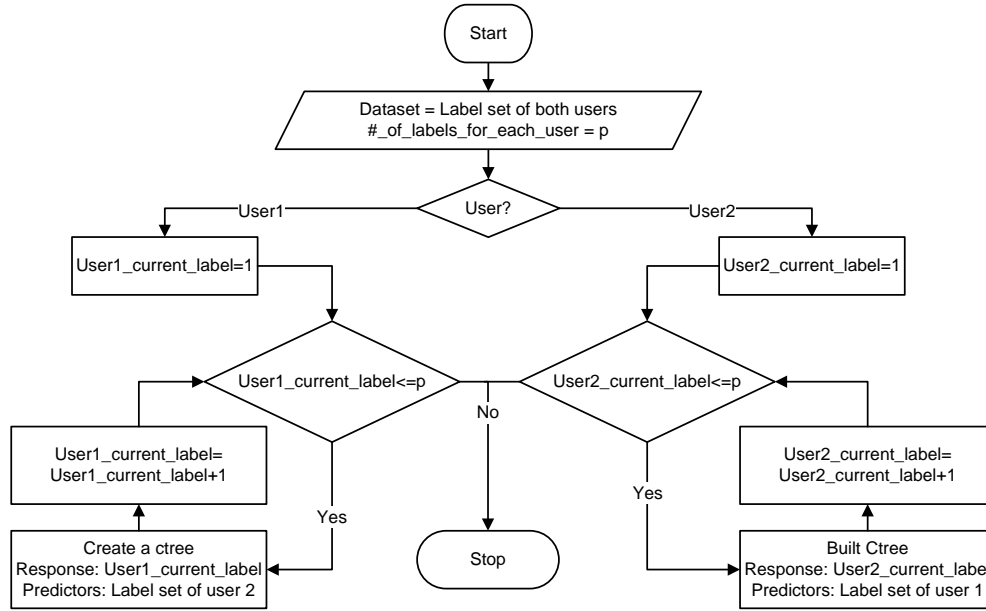


Figure 12. Identifying inter-user label dependency




Figure 9. Comparison between power-only and power-environmental models in terms of f-measure values for activities of *User 2*

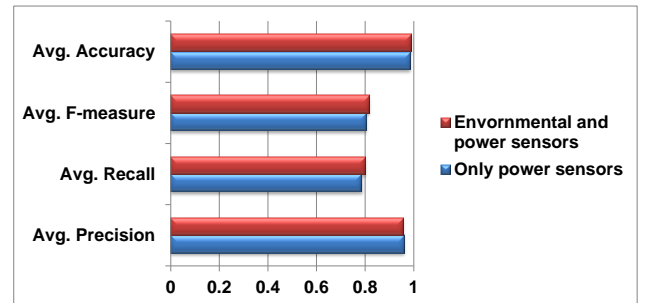


Figure 10. Comparison between power-only and power-environmental models in terms of average f-measure, accuracy, precision and recall

expose a significant positive or negative association relationship in the way they occur irrespective of the values of feature vectors. The first type of dependency we study focuses on dependency relations between the activities of an individual user and therefore uses individual datasets. The second type focuses on dependency relations between the activities of both users and therefore uses the combined dataset. These two different types can be termed as intra-user dependency and inter-user dependency respectively. We use the algorithm of conditional inference trees to study both types.

Figure 11 depicts the technique of identifying intra-user label dependency for each user. Initially, we define the number of labels for the respective user. For each individual label, we build a ctree model with this label as a response and the rest

of labels as predictors. This technique identifies for each label the set of labels that are statistically associated with it. This set identifies the dependency patterns for the respective label. This information helps understanding the occurrence of an activity given other activities of the same user.

Figure 12 depicts the technique of identifying inter-user label dependency. For this task, we use the combined dataset which contains the activities of both users. Our goal is to study the label dependency of each activity of *User 1* with respect to all activities of *User 2* and vice versa. The combined dataset consists of 22 labels. The first step is to decide which user's label dependency is to be studied. If *User 1* is selected, then for each label of *User 1* a ctree is built in which the response is the current label of *User 1* and the predictors are all the

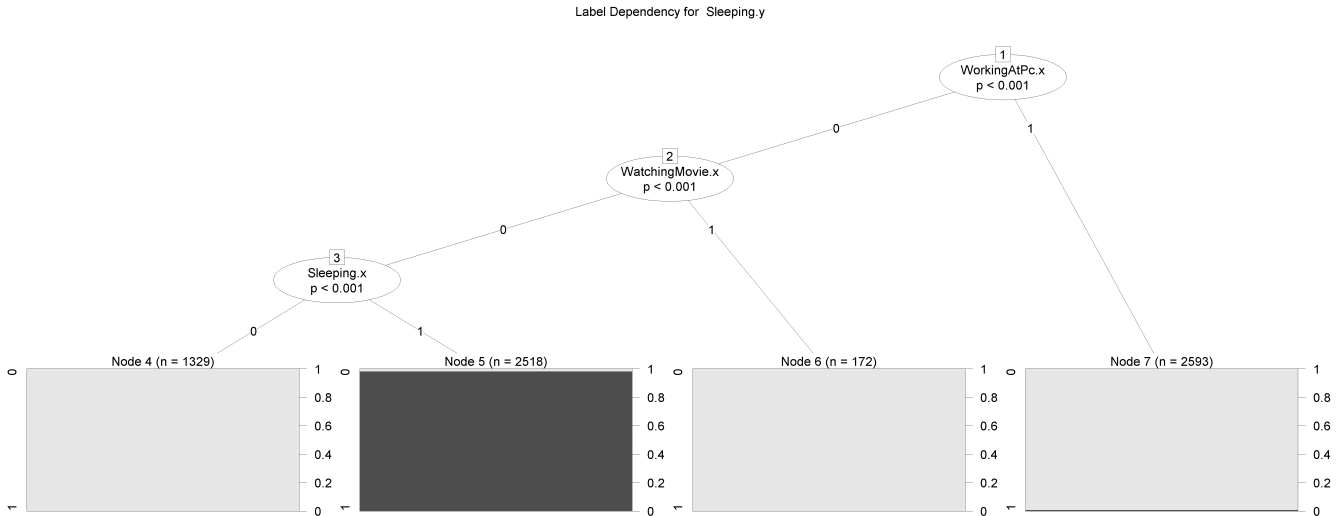


Figure 13. Combined dependency tree for *Sleeping* activity of *User 2* with respect to all activities of *User 1*

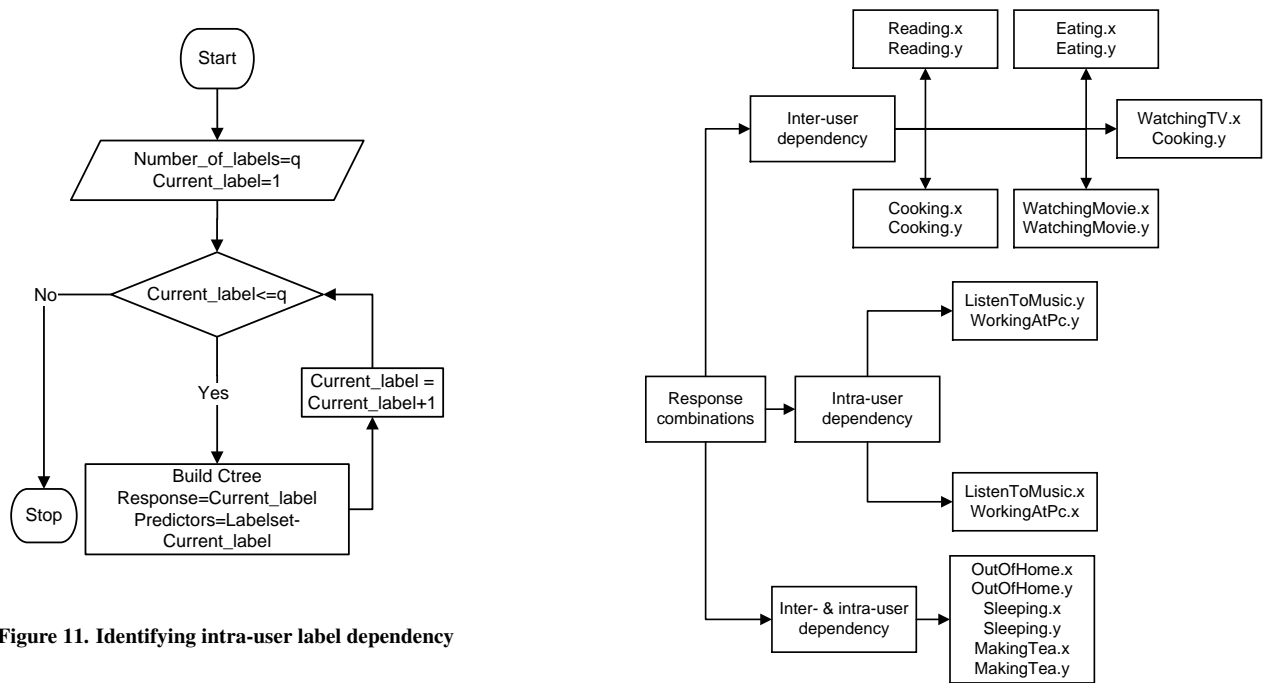


Figure 11. Identifying intra-user label dependency

Figure 14. Multivariate activity recognition model

labels of *User 2*. The same approach is repeated for the labels of *User 2*. This technique studies the co-occurrence of each user's activity with respect to all activities of the other user.

Figure 13 shows an example of inter-user label dependency for the activity *Sleeping* of *User 2*. As we can see from the figure, there is a significant positive correlation between this activity and the activity of *Sleeping* of *User 1*. Moreover, we notice a significant negative correlation between it and the activities of *WorkingAtPc* and *WatchingMovie* of *User 1*.

After building all ctrees for intra- and inter-user dependencies, we get the following information:

- Each of the activities *Sleeping*, *Reading*, *WatchingTV*, *Cooking*, *OutOfHome*, *Eating* and *WatchingMovie* is usually done by both users at same time.
- The activity *ListenToMusic* usually happens when the user is *WorkingAtPc*. This holds true for both users.
- The activities *OutOfHome*, *Sleeping* and *MakingTea* share a negative correlation i.e. they never happen together.




Figure 15. Comparison between dependency-based and power-only models in terms of f-measure values for activities of *User 1*




Figure 17. Comparison between dependency-based and power-only models in terms of average f-measure, accuracy, precision and recall




Figure 16. Comparison between dependency-based and power-only models in terms of f-measure values for activities of *User 2*

In order to utilize this information, we need to build a set of multivariate activity recognition models. Each of these models takes a set of dependent activities as its multivariate response variable. As the algorithm of conditional inference trees allows multivariate responses, we use it as a base classifier for building multivariate activity recognition models.

Based on the results obtained from intra- and inter-user dependency models, we combine correlated labels into a set of multivariate responses as shown in Figure 14. For each of these combinations, we build a ctree model with the combination of labels as a multivariate response. The predictor variables for these ctree models are the feature vectors comprising power sensors and temporal relations as shown in Eq. (3) after the exclusion of the values of environmental sensors. The final activity recognition model is therefore an ensemble of multivariate ctree classifiers. In the next section, we compare the performance of dependency-based model to the performance of power-only models. As explained in Section 4.3, power-only models are built using the approach of binary relevance and based on the feature vectors comprising power sensors and temporal relations between subsequent activities.

Evaluation

Figures 15 and 16 compare the predictive performance of dependency-based model to the performance of power-only

models. The comparison is done in terms of f-measure values achieved in recognizing activities of each user. As we can see from both figures, dependency-based model has significantly improved the predictive performance of activities *Reading* and *Cooking* for *User 1*. Moreover, it has increased the predictive performance of the activity *MakingTea* from 0% to 66% and 45% for *User 1* and *User 2* respectively. Figure 17 shows an overall comparison between dependency-based and power-only models. We notice from this figure that identifying and utilizing dependency relations between activities significantly improved the values of recall and f-measure by 12% and 11% respectively.

CONCLUSION

In this work, we presented our solution for activity recognition in multi-user environments. Our system is able to recognize activities of multiple users solely based on fine-granular measurements of power consumption and without the need for installing any other sensing modality. By conducting an extensive evaluation study, we proved the importance of dependency relations between co-temporal activities for enhancing the predictive performance of activity recognition models in multi-user environments. We identified, extracted and utilized these dependency relations for an individual user and between both users as well. Moreover, our evaluation showed that temporal relations between subsequent activities play an essential role in enhancing the predictive performance of multi-label activity recognition models.

ACKNOWLEDGMENTS

This work has been financially supported by the Social Link Project within the Loewe Program of Excellence in Research, Hessen, Germany and has been co-funded by the German Federal Ministry for Education and Research (BMBF) within the SMARTER project as well as by the German Research Foundation (DFG) in the framework of the Excellence Initiative, Darmstadt Graduate School of Excellence Energy Science and Engineering (GSC 1070).

REFERENCES

1. Rakesh Agrawal and Ramakrishnan Srikant. 1994. Fast Algorithms for Mining Association Rules in Large

- Databases. In *Proceedings of the 20th International Conference on Very Large Data Bases (VLDB '94)*. 487–499.
2. Alaa Alhamoud, Felix Ruettiger, Andreas Reinhardt, Frank Englert, Daniel Burgstahler, Doreen Böhnstedt, Christian Gottron, and Ralf Steinmetz. 2014. Smartenergy. kom: An Intelligent System for Energy Saving in Smart Home. In *Local Computer Networks Workshops (LCN Workshops), 2014 IEEE 39th Conference on*. IEEE, 685–692.
 3. Alaa Alhamoud, Pei Xu, Frank Englert, Andreas Reinhardt, Philipp Scholl, Doreen Boehnstedt, and Ralf Steinmetz. 2015. Extracting Human Behavior Patterns from Appliance-Level Power Consumption Data. In *European Conference on Wireless Sensor Networks*. Springer, 52–67.
 4. Rong Chen and Yu Tong. 2014. A Two-stage Method for Solving Multi-resident Activity Recognition in Smart Environments. *Entropy* 16 (2014), 2184–2203.
 5. Aaron Crandall and Diane J Cook. 2010. Learning Activity Models for Multiple Agents in a Smart Space. In *Handbook of Ambient Intelligence and Smart Environments*. Springer, 751–769.
 6. Krzysztof Dembczyński, Willem Waegeman, Weiwei Cheng, and Eyke Hüllermeier. 2012. On Label Dependence and Loss Minimization in Multi-Label Classification. *Machine Learning* 88 (2012), 5–45.
 7. Afsaneh Doryab and Julian Togelius. 2012. Activity Recognition in Collaborative Environments. In *The 2012 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
 8. Yu-Jin Hong, Ig-Jae Kim, Sang Chul Ahn, and Hyoung-Gon Kim. 2008. Activity Recognition Using Wearable Sensors for Elder Care. In *2008 Second International Conference on Future Generation Communication and Networking*, Vol. 2. IEEE, 302–305.
 9. Torsten Hothorn, Kurt Hornik, and Achim Zeileis. 2006. Unbiased Recursive Partitioning: A Conditional Inference Framework. *JOURNAL OF COMPUTATIONAL AND GRAPHICAL STATISTICS* 15 (2006), 651–674.
 10. Tuan Anh Nguyen and Marco Aiello. 2013. Energy Intelligent Buildings Based on User Activity: A Survey. *Energy and Buildings* 56 (2013), 244 – 257.
 11. Emmanuel Munguia Tapia, Stephen S Intille, and Kent Larson. 2004. Activity Recognition in the Home Using Simple and Ubiquitous Sensors. In *International Conference on Pervasive Computing*. Springer, 158–175.
 12. Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. 2009. Mining Multi-Label Data. In *Data Mining and Knowledge Discovery Handbook*. Springer, 667–685.
 13. T.L.M. Van Kasteren, Gwenn Englebienne, and Ben J.A. Kröse. 2010. An Activity Monitoring System for Elderly Care Using Generative and Discriminative Models. *Personal and ubiquitous computing* 14, 6 (2010), 489–498.
 14. Tsu-Yu Wu, Chia-Chun Lian, and Jane Y. Hsu. 2007. Joint Recognition of Multiple Concurrent Activities using Factorial Conditional Random Fields. In *2007 AAAI Workshop on Plan, Activity, and Intent Recognition, Technical Report WS-07-09*. The AAAI Press, Menlo Park.