Improving Activity Recognition in Smart Environments with Ontological Modeling

Zachary Wemlinger(✉) and Lawrence Holder

School of Electrical Engineering and Computer Science,
Washington State University, Box 642752, Pullman, WA 99164, USA
{zewemli,holder}@wsu.edu

Abstract. The problem of activity recognition in smart environments has produced multiple divergent paths of research in an attempt to improve the usability and usefulness of smart environments. In this paper we merge these research paths by defining a method for mapping smart environment sensor activities into an ontologically defined semantic feature space. We show that by using this approach we are able to improve activity recognition by between 5–20 %.

Keywords: Activity recognition · Ontological modeling · Ontologies · Semantic Web

1 Introduction

Recent advances in pervasive computing technologies have enabled the exploration of intelligent environments as a means of providing daily activity monitoring and cognitive support for aging populations. In order to meet this potential the environment must be able to efficiently and accurately recognize the state of the resident and in what activity, or activities, the resident is engaged.

Unfortunately, activity recognition in a real-world smart environment is a challenging task, partly due to the sparsity of the available data and its highly skewed class distribution. Also, variability between individuals and environments can amplify the difficulty and expense of collecting and annotating data required to learn in a new environment. While some researchers have proposed expert systems based on semantic modeling as a method to avoid dependence on machine learning algorithms, these approaches carry expectations related to sensor utility and activity structure which often do not transfer well to the real-world.

The main contribution of this paper is a process for integrating semantic knowledge into the process of learning activity recognition models. We model five existing smart environments using Semantic Web technologies and existing ontologies related to smart environments and demonstrate that activity recognition can be reliably improved by 5–20 %, depending on what measure is used.

C. Bodine et al. (Eds.): ICOST 2014, LNCS 8456, pp. 129–137, 2015.
DOI: 10.1007/978-3-319-14424-5_14

2 Ontological Models of Smart Environments

Several researchers have defined ontologies for smart environments, each with a slightly different focus. Of these ontologies, DogOnt[1] [1] and COSE[2] [16] are both publicly available. To the best of our knowledge the ontologies in [2,15] are not generally available.

DogOnt provides a rich ontology for smart homes with an emphasis on facilitating device interoperability. As such it provides many concepts related to device capabilities, functionality and commands providing an API for smart environments. This focus makes DogOnt a good ontology for use by intelligent agents when controlling and communicating with devices in a smart environment. COSE, on the other hand, is a smaller ontology focused on modeling objects and sensors within a smart environment. In [16], the authors make clear that one of their design goals was to integrate COSE with a top level ontology, namely OpenCyc [10]. The authors argue that mapping into an upper-level ontology provides extra portability for models utilizing COSE and enables more integration with the wider Semantic Web.

For this work, we have chosen to use COSE due to its richer model of objects in the environment. It is worth noting that, given the structure of the Semantic Web, using one ontology in no way prohibits the use of a different ontology when the need to express different concepts arises. Thus, applications in smart environments can easily reference concepts from both COSE and DogOnt whenever needed.

3 Smart Environments

In this section we discuss the details of the environments which we have modeled for this paper. For a more general discussion of smart environments, refer to [4].

We have modeled five environments using the COSE ontology. One, named Kyoto, is a testbed for smart home research and is used in seven out of the eleven datasets in this study. The other four environments are homes which have been instrumented with sensors in order to gather longitudinal data on activities, the names for these environments are Aruba, Cairo, Milan, and Tulum. Table 1 provides details about these environments.

Figure 1 shows the locations of sensors in the Kyoto environment. This environment is a three-bedroom apartment with two levels. The upper level contains the bedrooms and a bathroom, while the lower floor contains the living room, kitchen and dining area. Controlled experiments in this environment were conducted on the lower level of this apartment. Those experiments account for five of the seven datasets gathered in Kyoto.

[1] http://elite.polito.it/ontologies/dogont.owl.
[2] http://casas.wsu.edu/owl/cose.owl.

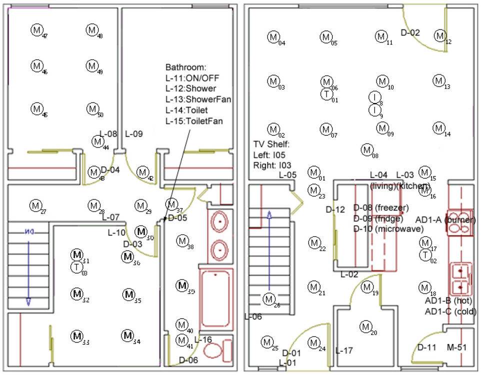

Fig. 1. Sensor layout for the Kyoto Smart Environment. This environment contains: 52 Motion sensors; 17 Light-switch sensors; 15 Door sensors; 12 Object contact sensors; 2 Temperature sensors; 2 Water flow sensors; 1 Home energy usage sensor; 1 Range burner sensor

Table 1. Environments used in this study

Name	Aruba	Cairo	Milan	Tulum	Kyoto
Sensors	39	32	33	18	102
Sensor Events	1,709,866	724,738	432,416	1,085,026	5,078,005
Sensor Types	Door, Motion, Temperature	Motion, Temperature	Door, Motion, Temperature	Motion	Door, Motion, Object, Power Usage, Temperature, Water Flow
Number annotated activities	9	9	7	13	118
Relevant Publications	[3]		[6]	[3]	[6,7,12,14]

4 Activity Recognition

Activity recognition is the task of recognizing when a person is performing a certain task. The set of possible tasks is unbounded, so smart home researchers generally consider a small number of tasks know as the Activities of Daily Living, or ADLs [9], which are of particular interest in elder care applications. These activities include: Grooming, Eating, Toileting, Bathing, and Personal Hygiene. The authors of [11] also suggest Instrumental Activities of Daily Living, or IADLs, which include: Using a Telephone, Shopping, Cooking, Housekeeping, Laundry, Taking Medications and Handling Finances.

These activities are the baseline around which much of the research, particularly health-care related smart environment research, has focused. Research groups tend to instrument an environment, observe research participants performing these activities, develop recognition algorithms, and then use the data to assess algorithm effectiveness.

4.1 Approaches to Activity Recognition

When designing systems to recognize activities, research generally falls into two areas. The first uses statistical machine learning techniques to learn models given observed data. The second uses logical models of activities defined *a priori* to build rules to determine what activity is being performed. Each method has benefits and drawbacks, though the machine learning approach is generally more popular.

The data-driven approach to activity recognition relies on instrumenting an environment with a range of sensors and using machine learning algorithms to mine the output of those sensors for patterns related to activities of interest. This approach is generally flexible and powerful enough to build useful activity models, though skewed class distributions and the need for data annotation make learning in new environments difficult.

In contrast, a "knowledge-based" approach is based on expert systems which contain activity models as a set of logical constraints. These models are used by inferencing engines to determine what activity is being performed. In [2], the authors propose such a system using Semantic Web technologies and a novel decision algorithm based on lattice-theory. The challenges when applying this approach are that getting data out of an environment which is clean enough to fit into a rule-based system is difficult. Also, building and extending these systems is difficult and costly.

In this work we demonstrate a hybrid approach which combines the best of both worlds in order to minimize the drawbacks of each while capitalizing on their strengths.

5 Learning with Semantics

The proposal is simple: Map sensor data into an ontologically defined feature-space for use by machine learning algorithms. This allows learned models to be applied in any environment where these mappings are defined, and if needed, integrate semantic rules into the activity recognition process.

5.1 Defining Semantic Space

Statistical machine learning algorithms necessarily learn by example. Given a dataset, the attributes for each example are considered to be the "feature space" of that dataset. When we say a "semantic feature space" we mean a feature space which is defined by concepts in an ontology. The challenge is to choose which concepts to use in our feature space.

To do this, we have taken inspiration from natural language processing techniques and have adopted an n-gram approach. With this approach, we define the feature space to be a set of n-grams built up from the concepts in COSE. In this approach each sensor is mapped to one or more n-grams based on the following.

First, let $\mathcal{S}$ be the set of concepts in COSE which are directly instantiatable in a smart environment. This excludes abstract concepts such as *Sensor*, but does include concepts such as MotionDetector. Next, for each sensor e of type t we find the set of concepts $C \subset \mathcal{S}$ to which e has a relationship. Further, let $\mathcal{B}$ be a set of predefined base concepts around which n-grams will be built. Consider $\mathcal{P}$ to be a partitioning of $\mathcal{C}$ into $|\mathcal{B}| + 1$ sets such that all concepts in set i are a specialization of the concept B_i. The last set in $\mathcal{P}$ contains all concepts in $\mathcal{C}$ which have no parent in $\mathcal{B}$. The features in semantic space to which e maps are its type t and the set of features in the cross product of the sets in $\mathcal{P}$.

As an example, sensor M014 is a motion detector which is situated above the table in the dining room in the Kyoto smart environment. For this work we let $\mathcal{B}$ be $\{Sensor, SpaceInAHOC\}$[3]. Thus, M014 is mapped to features that are the cross-product of the sets $\{MotionDetector\}$, $\{DiningRoom\}$ and $\{Table\}$. These features can be expressed with statements such as "Motion in the dining room" and "Motion in the dining room above the table".

5.2 Creating Feature Vectors

When creating feature vectors from streaming events we maintain the state of the environment as a vector $\boldsymbol{V}$. We allow sensor events to change this vector and every g time units we shift this vector onto a stack of state vectors which extend back for a limited amount of time; we call this $\boldsymbol{H}$ or the history matrix. Given a list of window sizes $\boldsymbol{W_s}$, we create a set of matrices $\boldsymbol{W_m}$ over $\boldsymbol{H}$ which define windows extending back a specific amount of time from the current moment. If we consider the state vectors to be the rows of the matrix, then the columns provide a history of each individual sensor.

The output feature vectors are, for each window, the mean of the columns plus the union of 1D FFTs run over each column. The real-valued inputs to the FFT mean that the result is symmetric and we only need to retain half of the FFT which helps to reduce excess noisy features. In order to map into semantic space, we observe that each ontological concept can be thought of as a sensor with a real value. The state of the sensor is simply the mean of the real sensors which map into the concept.

6 Experiments

In order to test our hypothesis that adding semantic models can improve activity recognition in smart environments, we have testing activity recognition algorithms over eleven datasets using three feature spaces.

6.1 Feature Spaces

The first feature space, which we refer to as sensor space, is a standard approach which directly takes features from sensor activities. In this space, each sensor produces two basic features for each time window: the sensor's mean activity and the dominant frequency of the sensor.

[3] SpaceInAHOC is short for Space In A Human Occupation Construct.

The second space is the COSE space, which is built by mapping sensor activities into semantic space using the procedure described in Sect. 5.1. The features in this space are the mean activity and dominant frequency of each n-gram.

Finally, we also have tested the union of these two spaces, which we refer to as the hybrid space.

6.2 Experimental Setup

Here we use window sizes of 3, 7, 14, 30 seconds and g is 0.5 seconds. As suggested in [3], we sample one feature vector every 10 seconds. All features are discretized using equal width binning with 5 bins then mapped to binary features. Feature vectors are labeled using the set of all labels observed on a sensor event during the 10 second window.

Activity recognition is inherently a multi-label classification problem and this is addressed here by training a single classifier for each activity in each dataset over each feature space and evaluating the performance of each model independently.

We have split each activity into n separate and temporally continuous sections. We train and test our models iteratively, such that on iteration i we train our model on the first i sections and test on the remaining $n - i$ sections. In our tests we have set $n = 10$.

Results are based on using Sofial-ML[4] which is a support vector machine implementation and which is partly described in [13]. This particular implementation allows calculating support vectors by stochastically selecting examples from both the positive and negative class. Doing this is necessary for our datasets due to the highly skewed class distribution where the median density of positive examples is only 2.8 %. Throughout these analyses, statistical significance is tested using student t-tests at the $p = 0.05$ level.

6.3 Results

In Fig. 2 we present performance using four metrics: accuracy, precision, recall, and RMSE. The RMSE is calculated on the error in the probabilities produced by our classifier. Due to the significant differences between precision and recall, we do not present the F-score directly. If the top ranked feature space has a statistically significantly improvement over other spaces it is marked with a diamond.

As the reader can observe, in 70 % of the cases, using a hybrid feature space provides a significant improvement over other feature spaces. In terms of performance relative to a standard sensor approach, using a hybrid space provides 5.1 % more accuracy, 20 % better precision, 5.6 % higher recall, and 9.7 % lower RMSE.

[4] Available at: https://code.google.com/p/sofia-ml/.

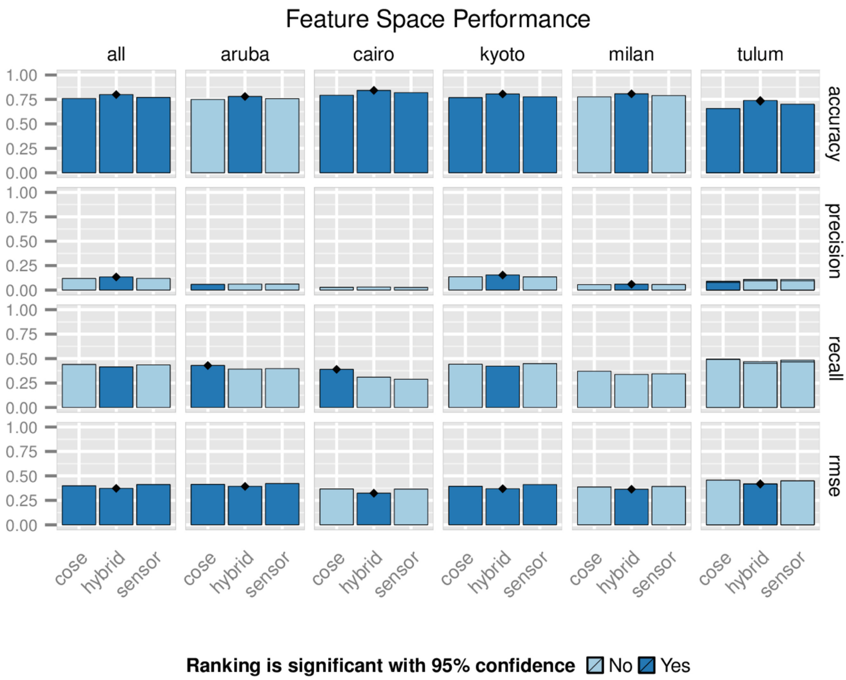

Fig. 2. Means for performance measurements given learning in the three feature spaces. Marked bars indicate the top ranked space with statistical significance. Significance is based on paired t-tests across all tests and activities in the environment. Note, RMSE is an error measure, so lower is better.

The sensor space used here is consistent with the feature spaces used in other data-driven approaches to activity recognition, e.g. [3,8]. Specifically, [3] evaluated performance of learning algorithms using several of the datasets utilized in this paper and reported an accuracy of 75 % using a hidden Markov model. The overall accuracy for the sensor space used here was 77 % and the accuracy for the hybrid space was 80 %.

7 Related Work

In Sect. 2 we discuss related ontological modeling efforts and in Sect. 4.1 we discuss ongoing research into activity recognition. In [3], Cook uses non-semantic method for mapping sensors into a common feature space with good effect. The authors of [5] provide an overview of transfer learning, which is highly related to this work. Transfer learning can avoid the need to create an ontology in order to map sensor events between environments, however in doing so also does not provide a method for integrating other semantic knowledge into the activity recognition process.

This work relates to these other efforts in that it provides a bridge between semantic modeling efforts and data-driven machine learning techniques. Utilizing the strengths from both of these areas of research holds promise for creating extensible and portable activity recognition systems.

8 Conclusions and Future Work

In this paper we have provided a method for integrating semantic knowledge bases into the activity recognition process for smart environments and have shown that this process provides a statistically significant improvement of 5–20 % to existing activity recognition approaches across a variety of environments and datasets.

While it is evident that logical rules can be applied when using a semantic feature space, we have not yet tested how effective such a system would be; doing so would be an immediate next step to this research. Other directions for future research include extending the work in [3] to learn novel activities using concepts from the ontology. Also, embedding natural language concepts into the ontology could provide the basis for intelligent natural language prompting systems to enhance smart home interactivity.

Acknowledgements. Thanks to the CASAS project at Washington State University for making the data used in this study available. This work is supported in part by National Science Foundation grant DGE-0900781.

References

1. Bonino, D., Corno, F.: DogOnt - ontology modeling for intelligent domotic environments. In: Sheth, A.P., Staab, S., Dean, M., Paolucci, M., Maynard, D., Finin, T., Thirunarayan, K. (eds.) ISWC 2008. LNCS, vol. 5318, pp. 790–803. Springer, Heidelberg (2008)
2. Chen, L., Nugent, C.: Ontology-based activity recognition in intelligent pervasive environments. Int. J. Web Inf. Syst. **5**(4), 410–430 (2009)
3. Cook, D.: Learning setting-generalized activity models for smart spaces. IEEE Intell. Syst. **27**(1), 32–38 (2010)
4. Cook, D., Das, S.: Smart Environments: Technology, Protocols and Applications, vol. 43. Wiley, New York (2004)
5. Cook, D., Feuz, K.D., Krishnan, N.C.: Transfer learning for activity recognition: A survey. Knowl. Inf. Syst. **36**(3), 537–556 (2013)
6. Cook, D.J., Schmitter-Edgecombe, M., et al.: Assessing the quality of activities in a smart environment. Methods Inf. Med. **48**(5), 480 (2009)
7. Dernbach, S., Das, B., Krishnan, N.C., Thomas, B.L., Cook, D.J.: Simple and complex activity recognition through smart phones. In: 2012 8th International Conference on Intelligent Environments (IE), pp. 214–221. IEEE (2012)
8. Krishnan, N.C., Cook, D.J.: Activity recognition on streaming sensor data. Pervasive Mob. Comput. **10**, 138–154 (2014)
9. Lawton, M.P., Brody, E.M.: Assessment of older people: self-maintaining and instrumental activities of daily living. The Gerontologist **9**(3), 179–186 (1969)

10. Matuszek, C., Cabral, J., Witbrock, M.J., DeOliveira, J.: An introduction to the syntax and content of Cyc. In: AAAI Spring Symposium: Formalizing and Compiling Background Knowledge and Its Applications to Knowledge Representation and Question Answering, pp. 44–49. Citeseer (2006)
11. Rashidi, P., Cook, D.J., Holder, L.B., Schmitter-Edgecombe, M.: Discovering activities to recognize and track in a smart environment. IEEE Trans. Knowl. Data Eng. **23**(4), 527–539 (2011)
12. Sahaf, Y.: Comparing Sensor Modalities for Activity Recognition. Master's thesis, Washington State University (2011).
13. Sculley, D.: Large scale learning to rank. In: NIPS 2009 Workshop on Advances in Ranking, pp. 1–6 (2009)
14. Szewcyzk, S., Minor, B., Swedlove, B., Cook, D.: Annotating smart environment sensor data for activity learning. Technol. Health Care **17**(3), 161–169 (2009)
15. Wang, X.H., Gu, T., Zhang, D.Q., Pung, H.K.: An ontology-based context model in intelligent environments. In: Proceedings of Communication Networks and Distributed Systems Modeling and Simulation Conference, vol. 2004, pp. 270–275 (2004)
16. Wemlinger, Z., Holder, L.: The COSE ontology: bringing the semantic web to smart environments. In: Abdulrazak, B., Giroux, S., Bouchard, B., Pigot, H., Mokhtari, M. (eds.) ICOST 2011. LNCS, vol. 6719, pp. 205–209. Springer, Heidelberg (2011)