Activity Prediction in Ambient Technologies

Isaac Smolund
Division of Science and Mathematics
University of Minnesota, Morris
Morris, Minnesota, USA 56267
smolu004@morris.umn.edu

ABSTRACT

In this paper we examine an emerging subsection of research within the field of *ambient technology*, or technology that occupies a less intrusive role in our lives, while offering the same convenience as more traditional devices. Within this field we examine the problem of *activity prediction*, wherein technology is programmed with the intelligence to observe its user's patterns and predict what their needs will be without prompting. To understand the field of activity prediction as it stands today, we discuss the work of various researchers in this area. These works include systems for the recognition of activity, for developing a smart-home environment utilizing activity prediction, and for assisting people with handicaps through technology. We also examine a paper seeking to establish a baseline activity prediction model and system of evaluation for other work in the field.

Keywords

ambient information, ambient intelligence, activity prediction, activity recognition, context awareness

1. INTRODUCTION

With every passing year, the way we go about our lives becomes more intricately tied with computerized and internet-connected technology. Most often, modern technology utilizes touch screens, the more traditional mouse and keyboard, or other button-based interfaces to understand what we want from them. However, these methods of interacting with technology are not without their drawbacks – some users may find these interfaces intrusive or inconvenient, and other individuals might have limitations in their lives that keep them from interacting with a particular technology in the intended way. The field of *ambient technology* seeks to solve this problem by creating new technologies that incorporate alternate forms of interaction, such as voice commands or motion detection, to provide a less intrusive way for our technology to observe our commands.

The goal of an ambient technology is to reduce its interference into our lives, by reducing the amount of interaction needed to operate it. Even with ambient systems, however, some kind of input by the user is still required. What if a technology knew what we wanted from it before we even told it? Researchers in the field of activity prediction are trying to create just that, a technology with the ability to figure out what we want based on what it observes of our actions and patterns of behavior. Different approaches exist in this field, and while each follows the same basic principles, there has yet to emerge a single ubiquitous method of predicting human activity in ambient technologies and *ambient intelligence* (AmI) networks - i.e., a system of multiple ambient devices working in tandem. In this paper, we will examine some recent attempts to create a system for improving people's lives through activity prediction in ambient intelligence which, while certainly notable and interesting in their own right, suffer from clear drawbacks in terms of documentation and testing. Because of this, we will also discuss a very recent attempt to create a single method for the creation and evaluation of activity prediction systems.

UMM CSci Senior Seminar Conference, December 2015 Morris, MN.

2. BACKGROUND

To understand the significance of activity prediction in this area, it is important to have a familiarity with the field of ambient technology. The goal of an ambient technology is to blend into our environment by functioning as an enhancement of an existing object, rather than adding yet another technological device into our daily routine. In a survey work [10], F. Sadri gives the following description of ambient intelligence:

> Ambient Intelligence is the vision of a future in which environments support the people inhabiting them. This envisaged environment is unobtrusive, interconnected, adaptable, dynamic, embedded, and intelligent. In this vision the traditional computer input and output media disappear. Instead processors and sensors are integrated in everyday objects. So for example, instead of using mice, screens, and keyboards, we may communicate directly with our clothes, household devices, and furniture, and these may communicate with each other and with other people's devices and furniture.

Clearly, ambient intelligence holds the potential to drastically change the types of technology in our life, and how we interact with it. However, there is still much work to be done in the field in order for ambient technology to become truly ubiquitous.

2.1 Advantages of activity prediction

While it's easy to see the advantages of having more technologically advanced versions of our everyday devices, the real determinant of their usefulness will be the software driving their behaviors. Any attempt to create a useful ambient technology will have to address the issue of how best to serve its users through the intelligence given to the objects involved.

The ability to accurately predict a user's actions before they make them would be an invaluable skill for any technology, and has been explored by researchers and developers for many years. For example, a 2011 paper by Phithakkitnukoon, Dantu, Claxton, and Eagle [8] details an attempt to increase the usability of personal cell phones, by predicting the calls a person will make and receive in a given day, creating a personal *intelligent address book* to give a person easy access to the numbers they are likely to call at a particular time of day in a particular situation.

2.2 Challenges

Three key challenges exist in the field of activity prediction. The direct prediction of a user's activities is the most obvious task for a technology in this field. Additionally, identifying and classifying a user's activities as they happen also represents a significant challenge. Finally, most activity prediction systems deal with self-evaluation in some way, in order to determine if a system is valid, and so that it may improve its predictions over time. While the majority of this paper will focus on the challenges of *prediction* and *evaluation* in these systems, here we will examine the challenge of *recognition*.

In order to predict the activities of a user, a system must be able to recognize their activities in the first place, in order to gather data to learn and predict from. Much of the work done in activity prediction assumes a recognition method is already available, and indeed several systems for human activity recognition have been developed.

One of these methods for activity recognition is detailed in Riboni and Bettini's paper, "COSAR: Hybrid Reasoning for Context-Aware Activity Recognition" [9]. The authors of this paper propose a system, known as COSAR (Combined Ontological/Statistical Activity Recognition), which attempts to accurately recognize human activity. Figure 1 shows a representation of the system architecture that COSAR relies upon, with sensors on a user's body giving information to a mobile device, which then feeds this information (along with observations from the device's own sensors) to a network server, where the majority of interpretation takes place, saving on efficiency by taking the computational load off of a user's device. Additionally, the COSAR system takes in information about the user's location to further refine recognition. For example, a user whose heartbeat is elevated, and who is currently located on a running trail is more likely to register as *'jogging'* than *'weightlifting'*.

Another aspect of activity recognition is dealing with anomalous data. Users with certain muscular disorders, for example, might move in a way other than what the system expects, or a user might otherwise diverge from the expected pattern of activity, for a number of reasons. The COSAR system attempts to deal with this eventuality by using a *historical variant* that consists of the system's sensor data for a short time just before a particular event. For example, if the system has recognized a particular segment of activity at five successive time stamps as $\overrightarrow{p} = \langle p_1 = jogging, p_2 = jogging, p_3 = brushingTeeth, p_4 = jogging, p_5 = jogging \rangle$, then p_3 likely represents an incorrect classification. The system would determine this by examining the p_1 and p_2 classifications, both of which were *'jogging'*. Since it is unlikely the user would perform such a drastic shift in activity in such a short time frame, COSAR would recognize that the reading of *'brushingTeeth'* was likely an error.

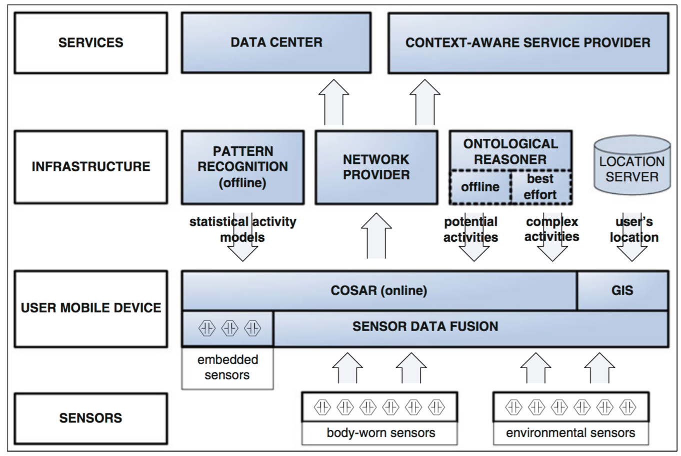

Figure 1: An illustration of the COSAR system. Taken from [9].

3. EXISTING SYSTEMS

In the attempt to find a useful, reliable way of predicting activity in ambient technologies, many researchers and software engineers have developed unique technologies to help their users in a variety of contexts. Here we will explore three activity prediction systems in two specific environments: A smart-home design, represented by the AdAPT system, and an assisted-care system for individuals with special needs, represented by the Hefestos and AmbienNet systems.

3.1 AdAPT

The AdAPT system, created by Jochen Frey and detailed in the paper "AdAPT– A Dynamic Approach for Activity Prediction and Tracking in Ambient Intelligence" [3], seeks to establish a system for the organization of smart technologies in an ambient intelligence-enabled home environment. Rather than fill each room with sensory devices to observe human activity, Frey's system attempts to recognize and predict behavior through a simpler sensor network, with more advanced intelligence behind it, using a system able to recognize the behavioral patterns of an individual, and use that recognition to better serve its user.

The AdAPT system's pattern-recognition abilities set it apart from many other systems. For instance, if the AdAPT system notices that a user turns on their AdAPT-connected radio at a certain time each day to a particular frequency, the system might begin to automatically switch on the radio to that station at that time each day. Another system, by contrast, might rely on a location sensor on the user's body, and turn on the radio when they approach it, or utilize voice commands to understand what it should do.

As we can see in Figure 2, the AdAPT system focuses on synthesizing different sources of data to form a more accurate idea of what a user's actions will be. To accomplish this, AdAPT takes in information not only from a limited sensor

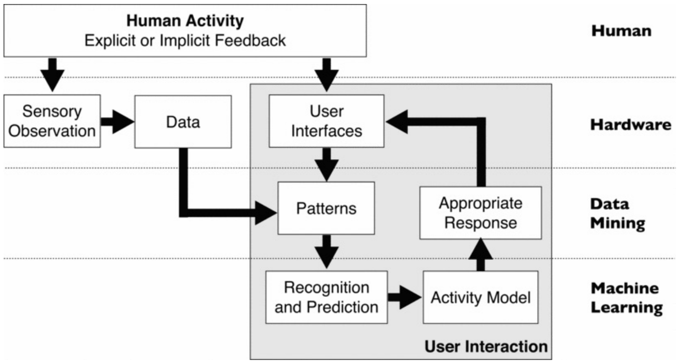

Figure 2: The interconnection of systems and human interaction in the AdAPT model. Taken from [3].

network in the user's environment, but also from direct interaction by the user, through an interactive display in both mobile phone and home television formats. The diagram divides this process into four layers - human, hardware, data mining, and machine learning. At the human layer, the system picks up information about the user and environment. This information is passed to the hardware level, where data is extracted from the sensory observations and, along with direct user input, is fed into the data mining and machine learning layers, which form a constant loop of interpreting user data and improving the system's knowledge about a user and their patterns.

Another unique feature of the AdAPT system is its cloud computing architecture. In his paper, Frey details the ways that his system utilizes cloud computing, and why this allows the system to better assist a potential user. Frey identifies four main ways in which cloud computing helps his system, namely *data storage*, *learning and adapting*, *recognition and prediction*, and *communication*.

Ideally in this system, a user's data would be stored securely in the cloud, allowing for more data to be stored with less intrusive hardware in the home. The system could also more accurately track and predict activity by learning human behavior not only from a single user, but from every user of the system simultaneously. Additionally, it could be given access to higher-level computational techniques, such as artificial neural networks, by utilizing the power of many machines working in tandem. Finally, by communicating between many systems which each hold information about their users, AdAPT can much better recognize patterns among many users of similar demographics or situations.

AdAPT's approach of using fewer sensor data than normal means that more interaction is required from the user, making the system less automated than others of its type. Though Frey mentions plans to preform a field test of the AdAPT system in comparison to other activity prediction systems, no results are discussed, making it difficult to say for sure how successful it would be in practice.

3.2 AmbienNet and Hefestos

Recent breakthroughs in computing have great potential for helping those in society who require special assistance, such as those who are disabled or elderly. Ambient intelligence is a particularly promising field for addressing this issue. In 2008, the authors of "AmbienNet: An Intelligent Environment to Support People with Disabilities and Elderly People' [6], proposed a system to assist with navigation and mobility of people who have difficulty walking and require a wheelchair for movement. What makes this system unique, however, is the amount of intelligent computing worked into its design.

The AmbienNet's "intelligent wheelchair" utilizes a wireless sensor network (using the Zigbee technology), along with ceiling cameras, to determine the location of the wheelchair relative to potential obstacles in a room. The system computes 'attractive' or 'repulsive' virtual forces of different objects in the room, based on what it believes the user is trying to accomplish. For example, the system might determine the user is attempting to navigate to a particular table, giving that object an 'attractive' force, while an obstacle on the ground, such as a toy or chair, would be given a 'repulsive' force. These forces, combined with the joystick input of the user, provide more accurate movement for users who might otherwise have difficulty navigating an obstacle-filled room.

Several years later, another group of researchers attempted to address the problem of how best to assist individuals with special needs. The Hefestos model, as detailed in "Hefestos: A Model for Ubiquitous Accessibility Support" [4], represents an attempt to assist people with disabilities or other impairments on a larger scale than AmbienNet. While the AmbienNet system attempts to help people who have difficulty with mobility, Hefestos seeks to address all different kinds of disabilities, and to this end relies heavily on a user profile system. By keeping records on the specific needs of individuals, the proponents of the system hope it will be able to better predict and cater to an individual's needs.

The Hefestos model uses four modules in its determination of how to proceed. The *User Profile module* mentioned above keeps track of specific data and preferences for a particular user. The *Special Needs module* contains information about a particular special need a user may have, and allows information about the user's activity to be standardized among all others with that special need. The *Contexts Module* looks to six contexts in a given situation - location, physical, temporal, social, special needs context, and accessibility - and allows the system to take these factors into account; for example, a public bus might only be expected to be available to a user in certain temporal and location contexts. The *Trails Module* keeps tracks of a user's contexts over a period of time, and is what allows the system to predict a user's future activity.

This model also proposes an ontology (a method of thinking about elements of a system and how they relate to each other) specific to the Hefestos model. This ontology describes the types of data the model keeps about people and their disabilities, and how the data points relate to one another.

The Hefestos model seeks to create a system of "ubiquitous assistance", and to determine its success in this task, the researchers behind Hefestos engineered a case study, detailed in their paper, to test the system. A quadriplegic student was given what the authors refer to as a *smart wheelchair*, an otherwise standard motorized wheelchair controlled by a smartphone running the Hefestos system. The student was then instructed to try out the system by exploring the campus of *Unisinos*, a Brazilian university. As the student navigated around the campus, the Hefestos system - which had internally marked the locations of various wheelchair accessible landmarks based on previous test runs - would give information about the nearby area based on its predictions of the student's behavior, drawing from contexts such as location, time, and the user's profile. For example, around midday, Hefestos suggested a cafeteria nearby the user's location with wheelchair-accessible seating. With further use, the system could continue to refine its predictive behavior based on the user's particular habits.

Feedback to the Hefestos system was generally positive, with the student describing the controls as "intuitive" and "easy to use" [4], however she noted the limitations of all user interaction taking place through a small cell phone screen. At the time of this writing, there has been one additional paper on the Hefestos system, published in 2013. In this updated paper, we can see the user interface is now much larger and clearer [5], and while this would likely be more helpful for users of the system, there is no data about additional field testing to reference. Additionally, though Hefestos strives to become a generic, ubiquitous model of accessibility, the system would do little to help someone with difficulty in an area other than mobility, such as visual or cognitive disabilities.

3.3 Evaluation of Activity Prediction Systems

While a sizable amount has been written about activity prediction, and many systems have been developed using activity prediction in ambient technologies, there are still some significant deficiencies in the field at this point. As can be seen in the papers we discuss above, there is very sparse documented testing of existing activity prediction systems, and often there is no evaluation present at all. This will certainly be an issue for the further development of these technologies, as it is very important that a system can be properly evaluated before it is accepted as useful and integrated into other work. However, in the time since the writing of the previously mentioned papers, there has been work done to create more reliable methods of development in activity prediction, and it is likely this trend will continue in the future.

4. IMPROVED EVALUATION METHODS

Since activity prediction in ambient intelligence is such a relatively recent field of computer science, consistent evaluation methods have yet to be developed to allow a reliable assessment of how successful and useful a particular approach might be for predicting user activity. A recent paper attempts to solve this problem by proposing a system of evaluation for activity prediction methods. In their 2015 work, "Data-Driven Activity Prediction: Algorithms, Evaluation Methodology, and Applications" [7], Minor, Doppa, and Cook give an overview of some of the different aspects and challenges of the activity prediction field, as well as proposing their own solution to the task. Their approach, however, is a much more adaptable one than those mentioned above, and could conceivably be applied to many activity prediction models. Perhaps most importantly, the researchers document a methodology for the reliable evaluation of activity prediction technologies, something that could be adapted by future research in the field in order to better determine the effectiveness of a particular system.

4.1 Learning Algorithms

The approach to the problem of activity prediction proposed by the writers of "Data-Driven Activity Prediction" is much more concrete and technical than previous published works. The writers find many advantages in their approach, and describe several accomplishments of their work: They *formulate* and *solve* the problem of activity prediction, and *reduce* it to a regression learning problem. Regression learning is a much more commonly addressed topic in the field of computer science research than activity prediction, and because of this the authors are able to much more effectively build off of previous bodies of work. Due to this, In this section, we detail the two approaches to activity prediction put forth by this paper under their philosophy of how the problems should be addressed.

The methods of activity prediction detailed in this paper assume that a system of activity *recognition* is already in place, and that user actions can be accurately recognized and encoded into data. The first algorithm described in the paper is called an *independent predictor* algorithm, or IP for short. This algorithm operates independently from much of the context information that other approaches use, such as data about time and relationships to other activities. The IP algorithm is simple, and fast to utilize and evaluate. However, it would not be as useful for making more complicated predictions, and is intended as merely a baseline algorithm for the paper's second algorithm.

The *recurrent activity predictor* algorithm builds upon the IP algorithm to be a fully-functional activity prediction system. While the IP algorithm can predict activity based upon specific moment-to-moment situations, the recurrent activity predictor algorithm incorporates context data to improve the accuracy and versatility of its predictions. The recurrent activity predictor is also able to interpret information about relationships between activities to further improve its predictions. In addition, the algorithm uses an *imitation learning approach*, in which a "learner" attempts to replicate the decision making of an "expert" in order to complete a task. The activity prediction implementation of this can be seen in Figure 3, where an overview is given of the basic steps the algorithm goes through to predict activity, as well as evaluate its own predictions. Pseudocode for the algorithm can be seen below in Algorithm 1.

Algorithm 1 RAP Learning Through Exact Imitation

Input: Λ = Training sequence of sensor event data labeled with activity segments, L = Loss function
Output: $\mathcal{F}$, the recurrent predictor

Initialize the set of regression examples D = $\emptyset$
for each time step $i = 1$ to $|\Lambda|$ **do**
  Compute local features $\Psi_{local}(i) = \Phi(\lambda_i)$
  Compute context features $\Psi_{context}(i)$
  Compute joint features $\Psi_i = \Psi_{local}(i) \oplus \Psi_{context}(i)$
  Compute best output $\mathbf{y}_i^* \epsilon \Re^T$ using the loss function
  Add regression example $(\Psi_i, \mathbf{y}_i^*)$ to $\mathcal{D}$
end for
$\mathcal{F}$ =**Multi-Output-Regression-Learner**($\mathcal{D}$)
return learned predictor $\mathcal{F}$

In the recurrent activity predictor system, a timeline of events is iterated through, and at each iteration (timestep, or i in the algorithm), multiple factors are calculated - the local features (data from the IP algorithm), the context features (information about timing and relationships), the joint features (a combination of local and context features), and the best activity predictions ($y_i^* \epsilon R^T$ in Figure 3), which are then added into a vector $\mathcal{D}$. This continues for as many elements there are in Λ, the vector containing the sensor data

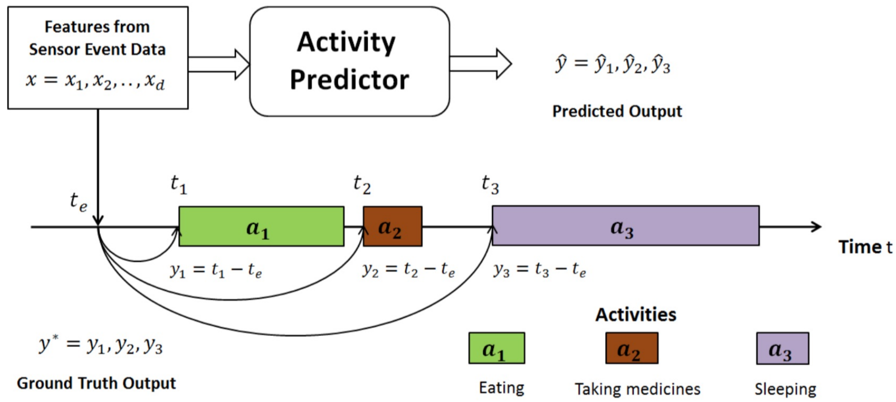

Figure 3: The method by which the "Data-Driven Activity Prediction" system evaluates its own predictions. Each t represents the time at which an event a occurs, in this case "Eating" (a_1), "Taking medicines" (a_2), and "Sleeping" (a_3). t_e is the time at which an initial event is observed which is known to lead to other events. The sensor data gathered at t_e is represented by the vector x, which is used by the *activity predictor* to generate $\hat{y}$, the *predicted output*, with each element of $\hat{y}$ representing the difference in time between t_e and the corresponding t. As the events occur, the actual difference in time is added to the vector y^*, the *ground truth output*, which is then compared against $\hat{y}$ to determine the accuracy of the predictions. Diagram from [7].

from each timestep. Following this loop, the *Multi-Output-Regression-Learner* $\mathcal{F}$ is calculated and returned. One weakness with this model identified by the authors of the paper is error propagation, i.e. if the algorithm imitates something incorrectly, this mistake could grow worse and worse as the system went on. A solution proposed for this is an existing system known as DAgger, which could be incorporated into the model in future development.

4.2 Evaluation Methodology

The writers of "Data-Driven Activity Prediction" introduce several metrics for evaluating the performance of an activity predicting system. For one, activity prediction can be evaluated by the *accuracy* of the system's predictions, and one can measure the similarity between the activity predicted by a system and the actual activity of the system's user. Additionally, such a system could be evaluated as what the paper calls a "forecasting algorithm", in which its accuracy is measured based on the difference in time of occurrence between the prediction and actual result. Minor et al. also acknowledge that any measurement of an activity prediction system is highly dependent on the system of activity recognition feeding into it, and that a system of activity prediction cannot be expected to perform well if given inaccurate data of a user's activities.

In the paper's evaluation metrics, $\hat{y}$ is used to denote a vector of predictions (in the form of the difference in time between a particular sensor event and successive events) and $\mathbf{y}^*$ represents a vector of true events corresponding with $\hat{y}$. T represents the total number of activities observed by the system for a given period. Using these, the paper gives several possible equations for finding the accuracy of a system. First, the *mean absolute error*, or MAE, can be used to measure the average error of a system's predictions:

$$\text{MAE} = \frac{\sum |\hat{y}_i - y_i^*|}{T}$$

Similarly, the *root mean squared error* (RMSE) performs a similar calculation, though the paper notes the squaring of terms would weight large errors more heavily than small ones:

$$\text{RMSE} = \sqrt{\frac{\sum (\hat{y}_i - y_i^*)^2}{T}}$$

To further improve the evaluation of activity prediction, the paper suggests using a method that considers the relative importance of each activity predicted. For example, a very important activity such as taking a medication would have to be predicted very accurately for the system to be useful, while something like buying groceries could happen in a reasonably large time frame could have more room for error. The *range-normalized RMSE* (NRMSE), defined below, would begin to address these concerns, with minimum and maximum functions incorporated into the calculation, allowing results to be compared from different data sets.

$$\text{NRMSE} = \frac{RMSE}{\max(y_i^*) - \min(y_i^*)}$$

Since the NRMSE does not have a clear normalization factor with which to determine the actual magnitude of errors, the authors then suggest a metric that normalizes each error value for a prediction by y_i^*, which represents the value of the activity to be predicted. The *mean absolute percentage error* (MAPE) can be seen below:

$$\text{MAPE} = \frac{\sum \frac{|\hat{y}_i - y_i^*|}{y_i^*}}{T}$$

This method has its own drawbacks, however, as at timestamps near the beginning of the event, where y_i^* is very low, small errors can have disproportionately large effects on the result of the equation. The means that a small number of outlier cases can unrealistically inflate the MAPE value.

To deal with the distorting effects of outliers, the authors introduce the *error threshold fraction* (ETF):

$$\text{ETF}(v) = \frac{\sum I(\hat{y}_i, y_i^*)}{T}$$

In this equation, $I(\hat{y}_i, y_i^*) = 1$ if $|\hat{y}_i - y_i^*| \leq v$, otherwise it is equal to zero. The numerator of the fraction represents the number of events with error below the value of v. ETF(0) will find the number of predictions with no error. If the ETF only approaches 1 for large values of v, it is a sign that a significant number of large-error outliers exist. Through a combination of the metrics put forth by this paper, future activity prediction solutions could likely be evaluated through much more reliable means than those detailed above.

5. FUTURE OF ACTIVITY PREDICTION

When we look at very recent papers on activity prediction in ambient technologies, a definite trend can be seen towards increased documentation and testing of new systems. In the paper "Significant Correlation Pattern Mining in Smart Homes" [1], written in 2015, Chen et al. propose a smart-home system to reduce CO_2 and greenhouse gas emissions and better conserve energy using data mining and activity prediction. Of particular interest, however, is the way in which the authors approach the more technical aspects of their work. The paper clearly documents several algorithms used in the system, as well as much of the process leading to their development. Additionally, the system has clearly been put through rigorous testing, and was compared to similar systems, with the results showing clear proof that the activity prediction and pattern recognition aspects of the authors' system provide a significant advantage over a system without them.

A shorter paper from 2015 also shows improvement over previous work. In "Automatic Sensor Data Stream Segmentation for Real-time Activity Prediction in Smart Spaces" [2], Cho et al. propose a system for assisted-living environments, targeting a similar demographic to AmbienNet and Hefestos. Similar to "Significant Correlation Pattern Mining", the authors of this paper evaluate their method in a very useful manner. While the authors do not directly draw from "Data-Driven Activity Prediction", their methods show clear similarities to the evaluation methods in that paper, which shows these methods becoming more commonplace in the field as a whole.

6. SUMMARY AND CONCLUSION

While the field of human activity prediction in ambient technology has inspired a fair amount of research in recent years, it is clear there is still much progress to be made before the results of this research can make a real impact on the kinds of technology the average person will utilize on a day-to-day basis. In this paper, a small amount of this research has been discussed, though we believe this to be fairly representative of the field as a whole at this point in time. While systems such as AdAPT, AmbienNet, and Hefestos hold the potential to make a difference in the lives of users around the world – especially those with various handicaps – the lack of true testing and evaluation means that there is no way of knowing how reliable they really are. However, the trend in activity prediction research of primarily speculative articles seems to be shifting.

In particular, the system detailed in "Data-Driven Activity Prediction: Algorithms, Evaluation Methodology, and Applications" is of note as an example of a body of work that could create a baseline activity prediction model for future work to follow. It seems likely that future work in the field of activity prediction in ambient intelligence will draw heavily from works like that of "Data-Driven Activity Prediction", and that this will lead to many changes for the better in the field as a whole.

7. ACKNOWLEDGMENTS

This writing would not have been possible without the guidance and support of many people, in particular KK Lamberty as a direct adviser to the paper, Nic McPhee and Elena Machkosova as professors of the Senior Seminar class, Sara Lahr as an alumni reviewer, and many of my peers in the UMM CSci division as emotional support.

8. REFERENCES

[1] Y.-C. Chen, W.-C. Peng, J.-L. Huang, and W.-C. Lee. Significant correlation pattern mining in smart homes. *ACM Trans. Intell. Syst. Technol.*, 6(3):35:1–35:23, Apr. 2015.

[2] H. Cho, J. An, I. Hong, and Y. Lee. Automatic sensor data stream segmentation for real-time activity prediction in smart spaces. In *Proceedings of the 2015 Workshop on IoT Challenges in Mobile and Industrial Systems*, IoT-Sys '15, pages 13–18, New York, NY, USA, 2015. ACM.

[3] J. Frey. AdAPT – a dynamic approach for activity prediction and tracking for ambient intelligence. In *Intelligent Environments (IE), 2013 9th International Conference on*, pages 254–257. ACM, 2013.

[4] C. C. A. Y. João Tavares, Jorge Barbosa and R. Real. Hefestos: a model for ubiquitous accessability support. In *Proceedings of the 5th International Conference on PErvasive Technologies Related to Assistive Environments*, PETRA '12, New York, NY, USA, 2012. ACM.

[5] C. C. A. Y. João Tavares, Jorge Barbosa and R. Real. A smart wheelchair based on ubiquitous computing. In *Proceedings of the 6th International Conference on PErvasive Technologies Related to Assistive Environments*, PETRA '13, New York, NY, USA, 2013. ACM.

[6] A. M. R. C. Julio Abascal, Borja Bonail and J. L. Sevillano. AmbienNet: an intelligent environment to support people with disabilities and elderly people. In *Proceedings of the 10th International ACM SIGACCESS conference on Computers and accessibility.*

[7] B. Minor, J. R. Doppa, and D. J. Cook. Data-driven activity prediction: Algorithms, evaluation methodology, and applications. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, pages 805–814, New York, NY, USA, 2015. ACM.

[8] S. Phithakkitnukoon, R. Dantu, R. Claxton, and N. Eagle. Behavior-based adaptive call predictor. *ACM Trans. Auton. Adapt. Syst.*, 6(3):21:1–21:28, Sept. 2011.

[9] D. Riboni and C. Bettini. Cosar: Hybrid reasoning for context-aware activity recognition. *Personal Ubiquitous Comput.*, 15(3):271–289, Mar. 2011.

[10] F. Sadri. Ambient intelligence: A survey. *ACM Comput. Surv.*, 43(4):36:1–36:66, Oct. 2011.