

ANNOT: Automated Electricity Data Annotation Using Wireless Sensor Networks

A. Schoofs*, A. Guerrieri†, D. T. Delaney*, G.M.P. O’Hare*, and A.G. Ruzzelli*

*CLARITY: Centre for Sensor Web Technologies

University College Dublin, Ireland

†Dept. of Electronics, Informatics and Systems (DEIS)

University of Calabria, Italy

Abstract—Recent advances in low-power wireless networking have enabled remote and nonintrusive access to households’ electric meter readings, allowing direct real-time feedback on electricity consumption to home owners and energy providers. Fine-grained electricity billing based on appliance power load monitoring has been investigated for more than two decades, but has not yet witnessed wide commercial acceptance. In this paper, we argue that the required human supervision for profiling and calibrating appliance load monitoring systems is a key reason preventing large-scale commercial roll-outs. We propose ANNOT, a system to automate electricity data annotation leveraging cheap wireless sensor nodes. Characteristic sensory stimuli captured by sensor nodes placed next to appliances are translated into appliance operating state and correlated to the electricity data, autonomously generating the annotation of electricity data with appliance activity. The system is able to facilitate the acquisition of appliance signatures, training data and validate the monitoring output. We validate the concept by integrating the automated annotation system to the RECAP appliance load monitoring system.

I. INTRODUCTION

A. Energy-Efficiency and Green Incentives

Electricity, representing 41% of the total energy used in the Unites States [1], has a major role in homes for heating, cooling, lighting and providing power to appliances. The \$3.4 billion investment from the US government to spur transition to smart energy grids—modernized electricity networks—is a tangible trigger in realizing the technology for an energy-efficient electrical system [2].

The most significant energy savings can be realized by changing people’s habits [3], [4]. Home owners are often unaware of how changes in occupant behavior, renovations, aging appliances, or newly installed appliances affect the energy performance of their house [11]. Consequently, informing effectively energy consumers enables behavior altering technologies.

B. Home Electricity Monitoring

The electricity drawn by electrical equipment is an integral part of the total energy used in residential buildings. Yet, home owners do not know what they are paying for when they pay their bills.

Electric meters are typically used to measure the power consumption of a house. Recent advances in low-power

wireless networking have enabled remote and nonintrusive access to households’ advanced electric meters—*smart meters*, facilitating both local and global control and management of electricity for utilities; real-time power consumption can be communicated to the local utility and energy provider for billing and monitoring.

Another class of devices—*home electricity monitors*, can be easily clipped onto the power cables of classical electric meters to physically capture the power signal and subsequently forward it to a PC-class processing unit for direct real-time feedback on electricity consumption to home owners, with the end-goal of heightening awareness as to cost accumulation (e.g. [5], [6]). Feedback from commercial products currently focuses on real-time watt per hour power consumption and translated price, and averaged consumptions over different periods. Such information is however incomplete for building a detailed overview of where energy is consumed.

Intelligent software systems that are able to provide fine-grained recognition of appliances in real-time, based upon a single home electricity monitor attached to the main switchboard, have been experimented on for more than two decades [8]. Yet, wide commercial acceptance has as yet not happened. This paper argues that a primary reason that prevents wide acceptance is the difficulty to profile and calibrate an appliance recognition system ad hoc to each home environment. In this paper, we propose a system using a temporary wireless sensor network (WSN) deployment to automate electricity data annotation and render any type of appliance recognition system free of human supervision at setup.

Section II will review existing appliance load monitoring techniques and associated methods for profiling and calibrating systems ad hoc to specific deployment environments, i.e. creation of appliance signatures, system training and system validation. This section will highlight the limitations of existing techniques and Section III will present the motivation of this work. In Section IV, we show that wireless sensor networks can be used for automating data annotation. We evaluate this technique in Section V, based on tests realized together with the RECAP appliance recognition system developed at CLARITY [7]. Section VI brings forward future works and Section VII concludes.

II. RELATED WORK

Since the early work by G. Hart at MIT in the early 80s [8], many techniques for appliance load monitoring have been proposed and the associated approaches and procedures are recalled here.

A. Approaches to appliance load monitoring

The common goal to all appliance load monitoring techniques is providing a fine-grained electricity decomposition per appliance. This is typically done via the identification of appliances' specific characteristics or *signature* within a congregated signal. Three principal techniques have been investigated for recognizing appliance signatures and deriving electricity usage per device:

- Power signal step matching;
- Power line noise recognition;
- Appliance operation ambient sensing.

These techniques discriminate appliances based on appliances' specific power load patterns, powerline noise and emitted sensory information.

The Nonintrusive Appliance Load Monitoring (NALM) approach developed by Hart [8] uses power loads retrieved from electricity monitors. It partitions the normalized power values into segments—steady and changing periods, to characterize the power signal in successive steps or events, and matches them to appliance signatures. This basic principle has been reused by others with additional considerations for improving the appliance recognition accuracy. For instance, the load disaggregation algorithm [11] differentiates edge cases with a classification of appliances according to their frequency of use to balance the decision making.

Patel et al. [12] detect the electrical noise on residential power lines created by the abrupt switching of electrical devices and the noise created by certain devices while in operation.

With ViridiScope, Kim et al. [13] propose an indirect sensing approach to evaluate the power consumption of virtually every appliance in the home. They combine ambient sensors placed near appliances, measuring the emitted sound and magnetic field variations when appliances are *on* or *off*, and a single electricity monitor retrieving the total power load. Sensors are an integral part of the whole appliance load monitoring (ALM) system and are used throughout the whole process.

B. Issues with current ALM system calibration

The aforementioned ALM techniques require a set-up phase for system calibration. System calibration is a three-step procedure:

- (a) Generation of power/noise/other appliance signatures;
- (b) System training;
- (c) System validation.

Current techniques for ALM system calibration require human supervision, which introduces certain issues:

1) *Signature acquisition*: Signature acquisition demands that ALM systems must know what appliances are *on* when a power signal or a powerline noise is observed. Users need to turn on and off appliances one by one and record the appliance signature. This effort is time-consuming and increases with the number of appliances to record.

2) *System training*: The generation of training data for an ALM system is centered around the system user. If say one day of annotated power data is needed for the system training, the user will need to annotate what appliance is *on* at any instance. ALM systems can also self-generate training data by producing combinations of the set of signatures, but the results are not as accurate since power patterns may be emergent and different from the simple addition of appliances power signatures. For instance, appliances connected to the same line as an appliance drawing high current may suffer voltage fluctuation, thereby generating an overall power signal lower than the one produced from adding each appliance power signature.

3) *System validation*: Once the ALM system is ready to monitor appliances' loads, its accuracy needs to be verified. Typically, human observation is the preferred way to validate a system's accuracy. If the system outputs that appliances A, B and C are *on* at a specific time, the user verifies the veracity via direct observation.

III. MOTIVATION

This work is motivated by the observation that mature research in the area of appliance load monitoring is as yet not transferred to actual commercial integration. Hart [8] proposed a new approach two decades ago, transitioning from complex hardware with simple software to simple hardware with complex software for signal processing and analysis. Although Appliance Load Monitoring (ALM) poses to be a powerful and inexpensive tool for disaggregating electrical costs for a domestic setting, its potential is yet to be realised in the commercial sphere.

Current ALM systems require complex calibration and verification to be carried out by a trained technician before it can be used in a domestic environment. This required setup stage proves very costly and is an inhibiting factor for a large scale roll-out. Other systems have facilitated the setup phase via a guided installation procedure, in order that any home owner can conduct ALM system calibration. However, the correctness of the technical and manual input from background-free users essential to the system accuracy cannot be assumed. A solution for automatic annotation used in conjunction with an ALM system can eliminate the need for in depth technical knowledge to carry out this setup stage allowing the end user to install and use the system with minimal effort. We believe by reducing the costs at this stage we will begin to harness the benefits of ALM for a domestic environment.

IV. ANNOT: A WIRELESS SENSOR NETWORK TO AUTOMATE ELECTRICITY DATA ANNOTATION

With ANNOT, our approach to reducing human supervision in ALM systems consists of deploying a temporary wireless sensor network with sensor nodes attached to each appliance, and using the generated sensor data to autonomously annotate data during system calibration, as depicted in Figure 1. ANNOT is in principle applicable to any kind of ALM system. In the following we develop further the concept and present how the various steps of ALM system calibration are carried out within ANNOT.

A. Concept

Data annotation plays a major role in ALM system calibration for the acquisition of appliance signatures and system training and validation. Annotation of the data retrieved from ALM devices, whether it be from an electricity monitor, a power line capturing device or a permanent wireless sensor network, provides the information of what appliances are *on* at any time.

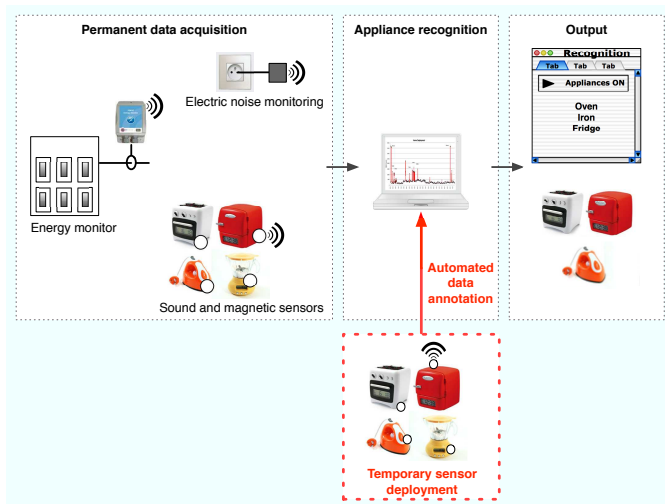


Fig. 1. Integration of ANNOT with appliance load monitoring systems. A temporary WSN deployment provides complementary information about appliances operating states.

1) *Sensory information for categorizing appliance activity:* ANNOT proposes the use of sensors for collecting information on the operating state of appliances to autonomously annotate electricity data. Multiple types of sensory data can be used to identify the operating state of an appliance:

- Temperature
- Light
- Sound
- Vibration
- Current variations

Using sensory information to detect appliance operating activity has been demonstrated with cheap sound sensors and indirect power sensing [13]; we show that using cheap temperature, accelerometers (vibration) and light sensors offer

alternatives. Jiang et al [14] designed ACme for monitoring AC electricity usage. We use the ACme node to include information from appliances' current variations in our study. Not all sensory information is helpful in categorizing an appliance operating state; we investigate in the following the appropriate combination of sensors to optimize cost, complexity and practicability. Furthermore, as we intend to only categorize appliance activity into activity states, we focus on sensor data variations. The latter consideration relaxes the constraints related to sensor calibration and the peculiarity of sensor data variation between appliances.




Fig. 2. Kettle temperature information as unique discriminating factor for understanding appliance operating activity.

2) *Single sensory information as discriminating factor:* Direct capture of current variations with the ACme node would in itself be sufficient as activity discriminator. Indeed, as shown in Figures 4–6, current is drawn when appliances are active. However, the associated cost to providing plug monitoring devices together with the difficulty of connecting them to many home appliances such as embedded ovens or dish-washers lower the advantages.

Ambient sensors offer a cheaper and more complete way of capturing the operating state of home appliances. Appliances generate a rich cocktail of heat, sound, light or vibrations when *on*. Using a minimal set of sensors per node reduces software size, data storage and communication, subsequently reducing cost and complexity. Using one single sensor per node i.e. per appliance would be the optimal approach, but one needs to understand the limitations. We illustrate how a temperature sensor can be used to identify the activity of an appliance emitting heat. Figure 2 shows a temperature measurement realized on a kettle after it being switched on. The delta in temperature increase, when the appliance is *on*, is a good indicator of an operating state change from *off/idle* to *on* state. Yet, as Figure 2 shows, the temperature does not directly follow the kettle operating state; temperature stays high long after the kettle is switched off. Generally, appliances emitting heat will keep on emitting heat for some time after being switched off. Subsequently, characterizing an appliance emitting heat as *off* or an appliance still emitting heat that is switch on again is not possible with a single temperature sensor. An automated data annotation system has to detect all state transitions, and this example shows that a combination of sensor information may be required. Furthermore, external events such as appliance movement or environmental condi-

tions should not lead to erroneous conclusions. Combining different sensory inputs reduces false positives.

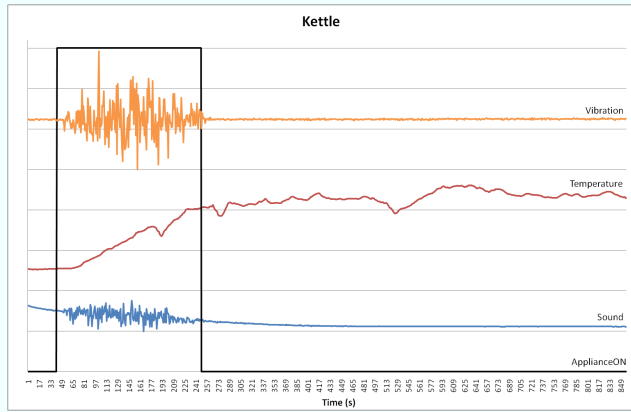


Fig. 3. Kettle multiple sensory information as discriminating factor for understanding appliance operating activity.

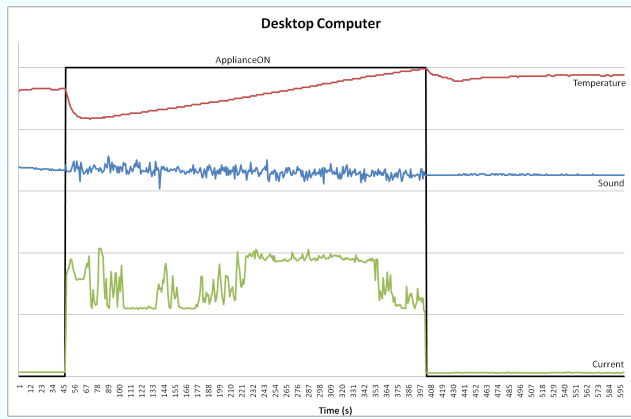


Fig. 4. Desktop computer sensory information exhibits operating activity.

3) *Combination of sensory information for accurate state recognition:* Types of sensory information emanating from home appliances are manifold. Figures 3–8 are examples of data that can be collected indirectly (via sensors) and directly (via capture of current draw) from appliances. This data has been scaled and translated to improve their observation. There is a clear correlation between appliances’ activity and sensory data patterns. For instance, an active fan can be identified from the vibrations generated, and whether the fan is rotating will be known from recurring light patterns due to nearby lamps or windows, see Figure 5.

B. System calibration steps

In this section, we describe how the different steps of ALM system calibration are facilitated with ANNOT and how it compares to state of the art techniques. We focus on ALM techniques based on power load as this is the method we are using for our research.

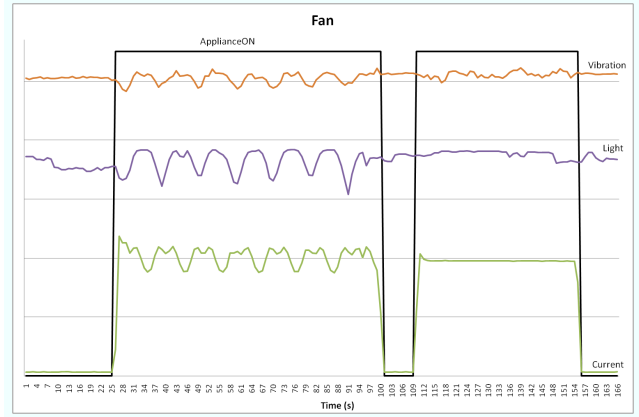


Fig. 5. Fan sensory information exhibits operating activity.

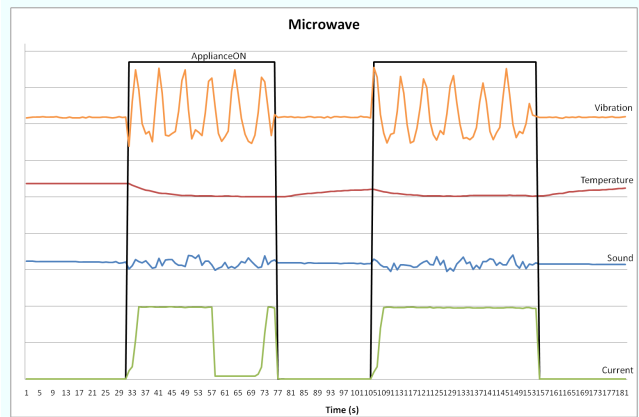


Fig. 6. Microwave sensory information exhibits operating activity.

1) *Obtaining appliance signatures:* Acquisition of accurate appliance load signatures is essential to achieving high-performance load monitoring. In addition to reducing manual intervention in the process, there are two desirable properties considered when generating appliance signatures:

- **Accuracy:** Generated signatures for an appliance should be as close to the load patterns observed by the system during operation.
- **Completeness:** Technique should obtain signatures encompassing all appliances in home.

ALM systems may be provided with signatures retrieved from previous deployments, where similar appliances were used. Signatures could be either already stored in the processing unit or accessed via remote connection to a signature database. Due to the immense number of existing electrical appliances and the fact that exact same appliances may draw different power, not considering household ad hoc appliance signatures violates the accuracy property and promises inaccurate predictions.

Another common technique is to successively plug an electricity monitor to each household appliance and retrieve the exact appliance load characteristic. The drawbacks of this technique are twofold: 1) the appliance signature captured at




Fig. 7. Monitor sensory information exhibits operating activity.

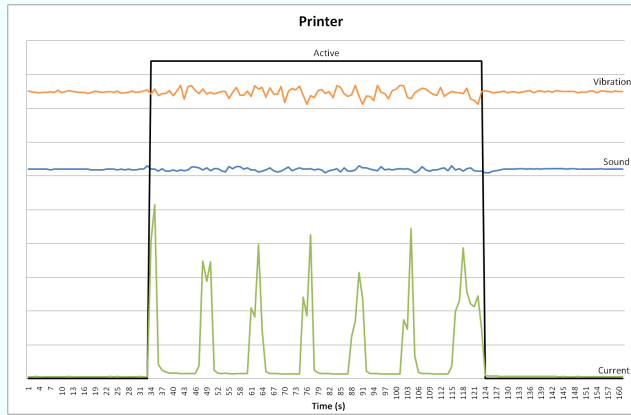


Fig. 8. Printer sensory information exhibits operating activity.

the appliance is different that the one seen at the electrical meter. This technique shows greatly improved accuracy but still violates the accuracy property, and 2) clipping electricity monitors to appliances' power cables is not always possible when cables are hidden or when appliances are of difficult access violating the completeness property.

Signature acquisition is then generally realized at the unit which processes the household power signal, so that learned load patterns are similar to those that will be observed when the system is running. The processing unit receives as input from the electricity monitor the total household energy consumed. The system can either try to identify recursive patterns or recognize patterns under human supervision. Both techniques require human supervision to indicate which appliance is *on* when certain patterns occur. This is generally done via a graphical user interface. With power line noise recognition techniques, a similar approach is used in collecting and recording noise signatures from appliances in the *on*, *off* and *idle* states [12]. Even though problems of variable power drawn by some appliances as well as concurrent *on/off* events affect this methodology, it is being used by most ALM systems as it provides a controlled and convenient way of associating load signatures to appliance activity for most of home appliances.

Using an automated annotation system to replace human supervision can fulfill both properties. As the signatures are generated by the ALM system, it inherently generates a signature that will be observed during operation. Using a WSN with multiple sensor types, and working on the assumption that every appliance generates auxiliary signals when active,

we can effectively use this technique for all appliances.

2) *Training and validation of pattern recognition algorithm*: The second step in ALM system calibration consists of training and validating the ALM algorithm so that it recognizes single appliances from multiple combinations of appliances. ALM system calibration follows three steps:

- **Generation of training data:** Produce power signal and annotate the observed appliance operational activity
- **Learning:** Input of training data into the learning algorithm
- **Validation:** Comparison of learning algorithm output and annotated training data

ALM systems typically self-generate training data from an initial set of signatures. Using automatic annotation, we are able to annotate complex loads over long periods. By storing this data in the system database, we supply a *ground truth* for the ALM system assumptions so it may calibrate itself accordingly e.g. measure the neural network weights.

V. CASE STUDY: EVALUATION OF ANNOT INTEGRATED TO THE RECAP APPLIANCE LOAD MONITORING SYSTEM


We performed an experimental evaluation to demonstrate the feasibility of automated data annotation in the context of appliance load monitoring. The experiment consists of evaluating whether data annotation can be fully automated and accurate, in order to demonstrate that real appliance load monitoring deployment and system setup phase can be realized without human supervision. The following describes the experimental setup, the technique used to correlate sensor data with electricity data, and depicts results of the experiment in a controlled environment.

A. Experimental setup


1) *RECAP system for appliance recognition*: The RECOgnition and Profiling of Appliances (RECAP) system is an intelligent system developed at CLARITY to recognise appliance activities in real-time [7]. The system benefits from recent advances on low-power wireless communication, where a PC-class processing unit receives data from a single electricity monitor clipped to the live wire of the electric fuse box. RECAP consists of two parts: (1) Electrical appliances profiling and generation of a database of unique appliance signatures, and (2) Artificial Neural Network (ANN) for appliances load recognition. The ALM system has achieved an average 87% recognition accuracy in deployment. Currently RECAP requires human supervision for generating appliances' signatures and observing the system accuracy.

The experiments use RECAP as the candidate ALM system to be made fully autonomous. The results of the experiments should be generalisable to any ALM system, as data annotation is independent of how power data is generated.

2) *Wireless Sensor Network*: We utilised a network of 25mm Tyndall motes [15] to capture sensory data next to home appliances. For our experiments, we equipped the motes with a 2.4GHz wireless transceiver and a multi-sensor layer including thermistor, humidity / temperature (digital), light



(a) Human supervised capture of appliances' signatures. Start/Stop events are manually notified.



(b) ANNOT autonomous capture of appliances' signatures.

Fig. 9. Illustration of how human supervision can be removed from ALM systems for generating appliances' signatures.

(LDR), sound, and accelerometer (3-axis). The nodes embed an ATMEGA 128L microprocessor on which we run TinyOS 2.1. We used the *Collection* protocol to retrieve data from the sensors over multiple hops when needed. Finally, we use the ACme node [14] to capture the current variations at the power plug. The ACme node uses the Analog Devices ADE7753 board to provide real, reactive, and apparent power calculations.

B. Architecture and data flow

The integrated system realizes data annotation following various steps as shown in Figure 10. Data is first sampled by the sensors and processed to identify activity. Next, the activity state is transferred to the ALM system controller and stored in the system database, for later retrieval during RECAP calibration. We develop in the following these steps and show results.




Fig. 10. Various steps in automated data annotation.

C. Data storage

Appliance state information sent from the sensor nodes is stored within the database used by RECAP. Figure 11 depicts portions of the database tables. A new *WSNApplianceData* table is created to store the appliance state information. One table is associated to each sensor node, and contains the appliance ID, the appliance state and a timestamp. Data will be available for query by the ALM system software.

D. RECAP calibration software modifications

The RECAP calibration software needs to be modified to incorporate sensory information in the calibration phase. We implemented a software program that combines electricity data and ANNOT data to create appliances' signatures, training data and data for system validation.

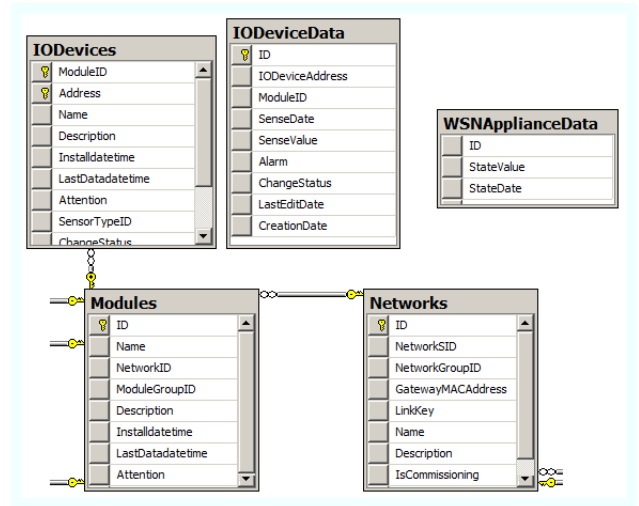


Fig. 11. Integration of appliance state information into the RECAP database.

1) *Generation of appliance signatures*: In RECAP, electricity data is retrieved from the electricity monitor and stored in the database. Like typical ALM systems, the system requires user input to signal when an appliance is turned on and off, so that the ALM system software can capture the electricity data that will form an appliance signature, see Figure 9(a). With ANNOT, user input is prevented by basing the capture of electricity for the signature on the input from the deployed sensors. In the example depicted in Figure 9(b), the kettle signature is recorded when the kettle activity is returned positive by the sensor node.

2) *Generation of training data and RECAP output validation*: Training data consists of generating electricity data annotated with appliance activity so that the RECAP system can train its neural network and refine its weights based on its output accuracy. Figure 12 depicts the procedure with a mock-up example. At any time, multiple appliances may be active, and the ANNOT system will provide that information. When the electricity data is processed by RECAP, the system knows what appliances are *on* for the input signal. The RECAP output or list of active appliances can then be compared to the input

list and validate the system accuracy.




Fig. 12. Generation of training data. At any moment in time, a list of active appliances is given by the wireless sensor network and can be associated to a power load.

E. Generation of appliance state from sensor data

Appliances' activity states provided to RECAP for electricity data annotation are generated by the sensor node attached to appliances. They are decided after processing based upon the combination of sensor outputs, as shown in Figure 13.




Fig. 13. Steps of appliance activity detection on sensor nodes.

1) *Detection of sensor data variation*: Temperature, sound, light and accelerometer data are sampled by the node. Data processing is realized on each sensor output to detect variations of sensor data.




Fig. 14. Appliance activity observable from variation in sensor data values.

TABLE I
DATA VARIATIONS OF SENSOR DATA IN IDLE STATE.

	Mean	Standard deviation	Min	Max	Max Delta (%)
Acc.	823.60	0.99	820.68	826.20	0.35
Sound	313.12	0.43	312	314	0.36
Light	205.66	5.01	158	211	23.17
Temp.	514.55	1.41	510	520	1.06

a) *Vibration and sound activity characterization*: Figure 14 depicts a typical output for accelerometer and sound data. Sensor data variations occur when the node is not in *idle* state. We use a sliding window approach and thresholds to detect this type of variation.

Data received by the microcontroller from the accelerometer and sound sensors is stored in a circular buffer, named x in this description, so that the variation detection algorithm reasons on the last set of input values. Our goal is to detect variations between successive data over a defined amount of time, as verified by Equation (a) below. Whether Equation (a) is verified, sensor data variation is either classified as negative or positive.

$$\sum_{i=1}^n |x[i] - x[i-1]| < nT \quad (a)$$

Equation (a) is the addition of the n successive absolute variations between the last $n+1$ input values. Conceptually, $x[n]$ is the last sensor input value. We compare this result with a threshold equal to nT , where T is the expected maximal variation of sensor signal values with respect to the signal mean value, in case of no sensor activity. With this comparison we want to establish whether consecutive variations are greater than a defined threshold, to conclude on sensor activity.

To attain a value for T , measurements were first carried out on the accelerometer and sound sensors before deployment. We sampled the sensors in *idle* state to measure the signal mean, standard deviation, min, max and maximal variation, the latter expressed as a percentage of the mean value. Results are shown in Table I. We see that the maximal variation—*Max Delta*—of the accelerometer data corresponds to 0.35% of the mean value, and 0.36% for the sound data. We concluded that configuring a value T equal to 0.5% of the acceleration data mean value and 0.4% of the sound data mean value for vibration and sound detection respectively would cover

all minor sensor variations when appliances are in *idle* state. These two values have subsequently been used to calculate the T values for processing the accelerometer and sensor data sets that we had captured. We also optimized the mean value calculation to be reactive to changes in average values when idle, re-calculating a new mean value based on the last set of points each time these points are found to be stable i.e. within the threshold.

We have experimented Equation (a) and the T value with the acceleration and sound data we have captured for the printer, shown in Figure 8, and with the desktop computer, shown in Figure 4, since their operation can be characterized with vibration and sound data. We chose to reason on the last four sensor values. For example, based on the a , b , c , and d points annotated on Figure 14, Equation (a) would translate at time $t = t_d$ into Equation (b), see below. In that specific case, Equation (b) would be verified and the sensor activity would be classified as negative.

$$|x[t_d] - x[t_c]| + |x[t_c] - x[t_b]| + |x[t_b] - x[t_a]| < 3T \quad (b)$$

Over time, the output is a series of sensor activity positive or negative classification, such as the one given in Figure 12. We then compared the results with the real sensor activity that we observed to compare the technique accuracy. Our results showed 94% accuracy for vibration detection and 93% for sound detection for the printer, and 94% accuracy for sound detection for the desktop computer.

b) Temperature and Light activity characterization: As observed on Figures 3, 4 and 5, temperature and light variations are not similar to acceleration and sound data variations. Variations are continuous and value differences between consecutive points are small. We compare temperature and light sensor outputs to ambient temperature and light respectively, to classify the sensor activity. Whether Equation (c) is verified, sensor data variation is either classified as negative or positive.

$$|x[i] - ambient| < T \quad (c)$$

Equation (c) compares the last value returned by sensors with the ambient value, and check whether the absolute difference is below a threshold T , where T is the expected maximal variation of sensor signal values with respect to the signal ambient value. Similar to acceleration and sound data, we attained a value for T by measuring the maximal variation of the temperature and light sensor outputs with respect to the mean value. Results are given in Table I; we configured a value T equal to 1.1% of the temperature ambient value for heat detection and 5% of the light ambient value for light detection. The ambient values are re-calculated when the appliances are in *idle* state to cope with changing day light and heat.

We experimented Equation (c) and the T value with the temperature data we have captured for the desktop computer, shown in Figure 4. Our results showed 92% accuracy in detecting temperature variation with respect to ambient temperature.

2) Appliance activity detection: Figure 13 depicts how the sensor variation outputs are combined to conclude on the appliance activity. Different combinations of sensor types are used to discriminate different appliances. Therefore, different *appliance filters* are needed to process the sensor variation outputs. For our experiments, we programmed nodes' software specific to each appliance, so that the programs knows what combination of sensor data should be considered when determining appliance activity. This hard-wiring approach has a main drawback of requiring specific code for each single node. We plan to investigate this matter for our in-home deployments. Sensor nodes will be equipped with an array of LEDs or 7-segment displays and buttons. Users will be given a look-up table where home appliances are associated to a binary code or an LED pattern. By pressing a button, home owners will easily set the LED pattern corresponding to the sensor appliance, and the software will know what appliance it is attached to, so that it can appropriately filter the sensor input. Subsequently, both sensor nodes hardware and software will be made generic.

We have achieved an 87% accuracy in detecting the printer activity, when combining the sensor variation output from accelerometer and sound sensors. We also achieved an 87% accuracy in detecting the desktop computer activity, when utilising the sound variation output and confirming it with the temperature variation. Temperature values do not immediately return to those in *idle* state when an appliance is switched off. We then use temperature data to confirm an appliance activity and do not use it as a discriminatory factor. A temperature variation alone would not be sufficient to characterize the computer activity as temperature still varies after the appliance is switched off. Using temperature variation reduces the accuracy that one would get with using only the sound variation detection i.e. 92% accuracy, but ensures that the detected activity is truly due to appliance activity and not to surrounding noise in that specific case.

These results obtained with simple processing shows that high accuracy can be achieved for detecting appliance activity, enabling accurate electricity data annotation and automation of ALM system setup. Furthermore, we showed that combining multiple sensor variation outputs slightly reduces the overall appliance detection accuracy but improves the confidence in the veracity of the appliance detection, which will be very important for home deployments where disturbing external factors may generate many false positives.

VI. FUTURE WORK

Immediate future work will aim at perfecting the appliance activity detection. Most of the errors generated at the sensor variation detection step were isolated and can be easily filtered out. Further extended pilots are planned with the RECAP system and the ANNOT system integrated. This will enable more extensive evaluation of system performance and usability.

VII. CONCLUSIONS

Data annotation plays a major role in appliance load monitoring systems. Currently, manual annotation is required and impedes large-scale roll-outs. We have presented ANNOT, a system that uses a network of wireless sensor nodes to automate electricity data annotation. Heat, vibration, sound and light are generated by appliances when they are active. We capture these sensory stimuli and orchestrate a load monitoring system calibration based on these inputs. We integrated ANNOT within the RECAP appliance load monitoring system and demonstrated the feasibility of such annotation technique with an 87% accuracy for detecting printer and desktop computer activity in our case study.

ACKNOWLEDGMENT

This work is supported by Science Foundation Ireland under grant 07/CE/I1147 and through the SFI-funded National Access Programme (NAP 252).

REFERENCES

- [1] Energy Information Administration, *Use of Energy in the United States Explained*, 2009
- [2] US Department of Energy, *President Obama Announces \$3.4 Billion Investment to Spur Transition to Smart Energy Grid*, 2009
- [3] A.H. McMakin, E. L. Malone, and R. E. Lundgren, *Motivating Residents to Conserve Energy Without Financial Incentives*, Environmental and Behavior, 2002
- [4] P.C. Stern, *Information, incentives and proenvironmental consumer behavior*, Journal of Consumer Policy 22, pp. 461–478, 1999
- [5] Ted: The Energy Detective, <http://www.theenergydetective.com>
- [6] P3 International Kill A Watt, <http://www.p3international.com/products/special/P4400/P4400-CE.html>
- [7] A.G. Ruzzelli, G.M.P. O'Hare, A. Schoofs, C. Nicolas, *Real-Time Recognition and Profiling of Appliances through a Single Electricity Sensor*, In Seventh Annual IEEE Communications Society Conference on Sensor, Mesh, and Ad Hoc Communications and Networks (SECON'10), 2010
- [8] G.W. Hart, *Non-intrusive Appliance Load Monitoring*, in IEEE, vol. 80, No 12, 1870–1891, 1992
- [9] Nonintrusive Appliance Load Monitoring Webpage, <http://www.georgehart.com/research/nalm.html>
- [10] Episensor Website, <http://www.episensor.com>
- [11] M. L. Marceau and R. Zmeureanu, *Nonintrusive load disaggregation computer program to estimate the energy consumption of major end uses in residential buildings*, In Energy Conversion and Management Journal, vol. 41, pp.1389-1403, 2000
- [12] S. Patel, T. Robertson, J.A. Kientz, M.S. Reynolds and G.D. Abowd, *At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line*, In Ubiquitous Computing Conference UbiComp 2007, 2007
- [13] Y. Kim, T. Schmid, Z. Charbiwala and M. B. Srivastava, *ViridiScope: design and implementation of a fine grained power monitoring system for homes*, In Ubiquitous Computing Conference UbiComp 2009, 2009
- [14] X. Jiang, S. Dawson-Haggerty, P. Dutta, and D. Culler, *Design and Implementation of a High-Fidelity AC Metering Network*, In the 8th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN09) Track on Sensor Platforms, Tools, and Design Methods (SPOTS 09), 2009
- [15] 25mm Tyndall Mote, <http://www.tyndall.ie/mai/Wireless%20Sensor-%20Networks%20Prototyping%20Platform.htm>