PRELIMINARY RESULTS FROM THE APPLICATION OF ADAPTIVE RESONANCE THEORY TO IDENTIFY THE OPERATIONAL STATUS OF SOLAR HOT WATER SYSTEMS BASED ON SMART METER WHOLE-HOUSE ENERGY RECORDS

David Menicucci
University of New Mexico
1 University Ave, ME Building
Albuquerque, NM 87106
dfmenic@unm.edu

Thomas P. Caudell
University of New Mexico
1 University Ave, EE Building
Albuquerque, NM 87106
tpc@ece.unm.edu

Anthony Menicucci
University of New Mexico
1 University Ave, ME Building
Albuquerque, NM 87106
amen@unm.edu

ABSTRACT

As a result of burgeoning installations of solar hot water systems (SHW), utilities have become concerned about the reliability of these systems and their real impact to the grid. If the systems fail or underperform, the utility must supply the shortfall of energy. Determining this information has been problematic for various reasons, one of which is the cost of field monitoring. A new idea has surfaced to use Adaptive Resonance Theory (ART) neural networks to analyze hourly whole-house energy data that are being collected by Smart Meters. ART is a computerized self-organizing learning network that mimics processing in biological neural systems. We applied ART to whole-house records of energy use that were recorded by Smart Meters. The results show that ART clearly identifies the date of installation of a SHW system. Similarly, ART was able to identify a simulated failure. Using these techniques we can measure the actual energy and power demand reduction from SHW systems without monitoring them directly.

Keywords: solar water heaters, smart meter, adaptive resonance theory, neural networks, reliability, energy reduction, power demand reduction, utility.

1. INTRODUCTION AND GOALS

Government incentives and utility rebate programs are spurring the installation of thousands of solar hot water (SHW) systems. The collective electricity reduction that results from these onsite residential generators has grown sufficiently large that some utilities have become concerned about their reliability and their actual installed energy performance. If these system fail to operate or if they fall short of their expected energy reduction goals, the utility has a legal duty to supply electricity to its customers to offset any shortfall. Furthermore, if the life and actual performance of SHW systems are unknown, then the utility cannot depend on their promised energy reductions and for planning purposes their potential value must be significantly discounted.

In 2008 Menicucci reported that in a field survey of 67 SHW systems in Florida, more than half had failed in the field (1). The same report also suggested that insufficient data exist to compute SHW systems' Mean Time To Failure (MTTF), a critical statistic that is needed by utility resource planners.

As a result a number of utilities are seeking to resolve questions centered in two areas of concern:

- *SHW energy savings performance*: What are the real energy savings from fielded SHW systems and how do those savings compare with the Solar Rating Certification Corporation (SRCC) energy savings ratings, which are used by some utilities to determine the size of installation rebates?

- *SHW reliability*: What is the reliability of SHW systems in the field and what is the mean time to failure for these systems?

This research effort addresses both questions by developing new and powerful analytical methods to analyze high frequency data from Smart Meters. These meters are currently used by many utilities and there is general agreement that they may hold the key to cost effectively monitor thousands of fielded SHW system and to generate large amounts of data that could be analyzed to determine the life cycle and grid-impact of SHW systems.

A primary goal in this research was to assess whether advanced data analysis techniques could effectively identify and measure the energy reduction signature of SHW systems that lay embedded in hourly Smart Meter residential data. If these signatures could be recognized and measured directly, the results could potentially provide the highest quality information to answer the two primary questions about SHW performance and reliability.

While such analysis has high prospective value, many analysts have eschewed it based on the assumption that the variance in whole-house data is so large that it would overwhelm and conceal the signature of the SHW system. Variance in energy use patterns is caused by changes in lifestyle (e.g., the addition of a new baby), alterations to the residence (e.g., a new major appliance), or weather, which influences the total home energy consumption.

However, a typical SHW system has an electric energy saving rating of around 3,000 kWh/yr, or about 20% of a typical annual all-electric residential energy load. Such a large signature should be recognizable, even with significant variance in the data.

A second goal was to determine if whole-house records could be used to assess whether failures of SHW systems could be identified in whole-house records, thus providing information about the mean time to failure of these systems. This reliability metric has eluded the SHW industry for nearly four decades.

2. TECHNICAL APPROACH

We believed that we could meet our objectives by analyzing whole-house Smart Meter data using Adaptive Resonance Theory (ART), at type of neural network system, along with standard statistical methods (2). Our idea was to use ART to analyze whole-house records to produce information that would allow us to identify the date when a SHW system was installed as well as the date of failure. Once we knew the installation date, we felt that we could isolate and measure the energy impact of the SHW system. If we knew the failure date, we could compute the Time To Failure.

Neural networks are machine-learning systems whose design and processing are inspired by biological nervous systems. They borrow from biology the notions of neurons as the elemental processing unit and the ways in which neurons are linked via unidirectional adaptive connections.

A network consists of a wiring diagram or graph showing the exact way a collection of neurons are interconnected. Some of the neurons have special roles in the network acting as interfaces to the environment; for example receiving inputs from sensors or sending signals to motor controllers. The connections in a network have the job of transporting the output of one neuron to the input of another, and are characterized by their source and target neurons as well as their connection strength, usually represented as a weight.

The processing of an individual neuron is as follows: First, the neuron integrates all of its weighted inputs arriving on in-coming connections from other neurons. Second it maps this integrated value through a possibly nonlinear function to form a new output. Third, it adapts the connection strengths on its in-coming connections. Finally, it broadcasts its new output to the network through its out-going connections.

The adaptive process is referred to as "learning" in this field, and falls into two broad classes: 1) supervised, and 2) self-organizing. The first method requires training data consisting of pairs of input/output samples. An input sample is supplied to the network through its input neurons, and the connection weights are modified to help the output neurons reproduce the output sample. The list of input/output samples is called the training set, and one learning pass through this set is called a training epoch. Through training over many epochs, the output neurons will gradually begin to match the desired output behavior provided by the training set.

The second method functions quite differently. In self-organizing learning, the network is not supplied a desired output for its output neurons. Instead, during the learning epochs, the weights are modified to help the output neurons autonomously encode categories of systematic or regular patterns that exist in the input samples. Often this is referred to as discovery learning, where the network learns to respond with a unique output pattern when a member of a category of similar input patterns is presented to it.

A neural network and its learning methods are referred to as a neural architecture. Neural architectures are typically implemented as algorithms in computer software simulations. The "goodness" of a learning method is usually quantified by the number of training epochs required to reach a given level of output performance.

ART neural architectures are in the class of self-organizing learning systems. When presented an input pattern, ART architectures will rapidly categorize it as a member of either an existing category or a new (novel) category. If an existing category matches the pattern, then the network will respond with an existing output code indicating its membership in a category. If no existing category matches, then the network will create a new output code that will in the future respond to the novel pattern.

We believed that ART could be used to identify the time of installation of a SHW system as well as the time of failure, if one existed in the time period of record under study. This would provide two benefits.

First, if many homes are studied and many failures are noted, then we could compute the mean time to failure of SHW systems.

Second, knowing the installation date with precision would allow us to compare the energy consumption and power demand characteristics from before and after the installation. If a sufficient number of homes were studied, then we could measure with fixed precision the energy and power demand reductions from SHW systems.

We believed that we could ameliorate the problem of high variance by applying moving averages to the data, which are effective in reducing noise and allowing subtle but systematic trends to be observed.

Our approach was to demonstrate the applicability of these methods by analyzing real data, i.e., hourly energy consumption records from Smart Metered homes.

3. DATA SELECTION

We obtained datasets of whole-house energy consumption records for seven homes. Each dataset contained a table of hourly energy consumption records for the home over a span of several years. We had no personal information about the home other than the street address and a digital identifier. But we did know for each home the date on which the utility issued a check for the installation rebate and the type of SHW system that was installed. Each dataset contained at least one year's data before and after the rebate date. We wanted to be reasonably sure that the installation date was included somewhere in the span of the dataset.

However, it is possible that the installation could have occurred weeks or even months before the rebate payout. A rebate is typically paid after installation and when all of the application paperwork is completed. The elapsed time for doing so could vary significantly among customers.

Our initial examination of the data involved a visual review of simple plots of energy use over time. We had to abandon two of the datasets outright because they contained large blocks of missing data that would have severely confounded the analysis and/or because they incorporated too limited a timeframe. Our analysis was performed on the five remaining datasets.

4. USE OF ART TO IDENTIFY THE INSTALLATION DATES

We applied ART processing to each of the five data sets with the goal to estimate the installation dates and to compare them with the payout dates. As discussed above, ART is a self-organizing system that learns in a manner similar to humans. It notices changes from patterns it experienced in the past. In our case we were seeking changes in energy use behavior.

As a first validation of ART's capability, we applied it to the raw data for each of the five homes in our database, in each case seeking an indication of the SHW installation. Since we were very certain that the installation must have occurred before the payout, any signal by ART of a later installation date would immediately challenge the veracity of the ART methodology.

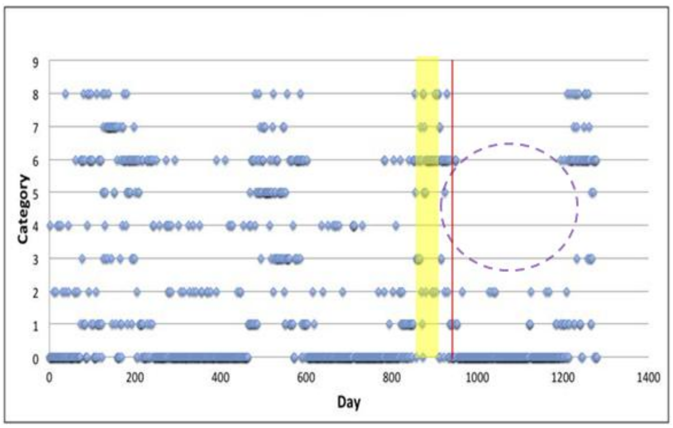

Fig. 1. Plot of one home's daily energy use category labels learned by the ART system versus day number. The day number corresponds to the date, with zero representing the earliest date in the file. Each blue dot represents a single 24 hour period.

The graph presented in Fig. 1 shows an example of the results of applying ART to one of the homes. Each category label represents a range of similar energy use patterns as determined by the ART system. A higher category label should not be assumed to be associated with a higher energy use behavior, although there may be some relationship.

In the plot the red vertical line represents the date of the rebate payout. The yellow rectangle represents our best judgment as to the range of dates when the actual installation occurred. That judgment is based on visually discernment of an alteration in the category patterns, the point at which the ART system began to identify a systematic change in energy use behavior. Simply stated, we expected that the installation would be exemplified by different groupings of the various categories that had been previously learned from the time before the installation.

The red dashed-line circle in Fig. 1 indicates the area of the plot that exemplifies the differences in energy behavior from similar periods in prior years. In this case ART's category 4 was essentially eliminated by an energy-related event that occurred subsequent to the time denoted by the yellow rectangle, which we presume to be the installation.

Effectively, ART did what a human would have done had he/she been familiar with the energy use behavior prior to the installation—it noticed a change in energy use activities.

The other four datasets showed noticeable changes in the energy behavior after the solar systems were installed, but they were not as profound as the ones shown in Fig. 1. However, the changes were sufficiently articulated to allow us to select the range of dates of actual installation.

In all five homes ART identified dates of installation that were either prior to the date of the rebate payout or around the same time as the payout, which is expected if the ART system were identifying the true installation dates.

5. TECHNICAL METRICS TO DESCRIBE THE SHW SIGNATURE

There are two important parts to the SHW signature: the energy savings and the power demand reduction.

Energy savings is a customary measure in the SHW industry and is defined as the difference between the following two energy metrics: 1) the amount of electricity that is annually consumed by a SHW system servicing a standard, 64 gal per day domestic water heating load, and 2) the electricity consumed by a standard electric water heater in the same location servicing the identical load. A SHW system is almost always configured with a backup system, usually a standard water heating tank that has been modified to supply heat when the solar system is not operating, such as at night. Therefore, even if a SHW system is operational it is nearly certain that some backup energy will be consumed.

We define the estimate of energy savings as:

$$Savings_{sol} = Q_{standard} - Q_{SolarBU} \qquad (1)$$

Where $Q_{standard}$ = the energy consumed by a standard water heater and $Q_{SolarBU}$= the energy consumed by a solar backup tank; both serving an identical load.

An estimation of the power demand reduction is based on the power rating of a SHW system. We estimate a SHW system's power rating by applying a method described by Menicucci in 2006 (3).

$$Power_{sol} = Savings_{sol} / (CF * 24 * 365) \qquad (2)$$

Where CF = the capacity factor for the solar system.[1]

This power rating is therefore analogous to the nameplate rating of a photovoltaic system or, in fact, any energy generator in the utility's system. For example, a SHW system with an energy savings rating of 3000kWh/yr would have a power rating of around 1.36kW. Effectively, when operating on an electric grid, this SHW system is a generator with a power rating of 1.36kW. Its energy product is not electricity, but rather heat that displaces electricity.

Demand, in utility terms, is a measure of the power requirements in a utility system and is an important consideration for resource planners. Power (expressed in kW) is a product of voltage and current. When a utility customer turns on an electric load, current flows from the grid into that load. The utility must supply sufficient current to maintain voltage above a legally defined lower limit. Voltage sag below the defined limit could damage customer equipment and create a liability for the utility.

To insure that enough current is supplied, the utility will install distribution hardware of adequate size to supply the largest expected current draw, even if that large event is rare. The extra equipment is expensive to install and maintain and those expenses are incorporated into the utility rates. Therefore, it is to everyone's advantage to minimize large current draws. The ideal load is one where there is virtually constant current flow, but that is rarely observed.

In an all-electric residence the large loads include the air conditioner, the space heater and the water heater. The water

[1] The capacity factor for most SHW systems in the US desert Southwest is around 0.25.

heater normally contains two electric elements, one in the upper third of the tank and one in the bottom third. On large draws of hot water, both elements in the water heater will be energized simultaneously. A common element is rated around 2kW producing a 4kW rating for the tank.

When the air conditioner turns on a large surge of current is needed to put the motors in motion. This initial locked-rotor motor draw is typically about six times the current that flows after the motor is turning. The current from a simultaneously energized water heater is added to that from the air conditioner and the maximum amount of current during the surge defines the power demand, even if the surge is brief.

The installation of a SHW system reduces demand in two ways. First, a solar hot water tank is typically a normal electric tank that has had the bottom electrical element removed. That bottom element is replaced by a solar loop heat exchanger. Thus, a solar tank would be expected to reduce demand by at least the rating of one electric element simply because only one element is installed in the tank.

Second, the solar tank's upper element would be expected to be energized less frequently than a normal tank's element because the solar collector is supplying heat to the tank. Therefore, the frequency of simultaneously operating the solar tank's element and the air conditioner or heater is lower than it would be with a standard tank. Simply speaking, there is less chance that the tank and other large loads will be operating together. In fact, at certain times of the year the solar tank's element may not be energized for weeks. When the effects of many SHW systems are aggregated on a feeder, the expected net result is lower overall demand.

Our database consisted of matrixes of hourly records of energy consumption from each of five homes. Each row in the matrix contained the energy record for the day. The energy consumed in each hour is presented in separate columns. Summary statistics for the day, such as the peak hourly energy consumption, are presented in other columns.

All of the values are in units of kWh. However, an hourly total kWh is equivalent to average power in kW over the hour, which is equivalent to *average* hourly demand. Thus, we can use the peak hourly energy value to directly represent the average maximum hourly demand for the day.

For example, on one day for one solar home the record might show total energy consumed was 26.5 kWh and the peak average hourly power demand was 2.4 kW.

In this particular example, suppose a solar system was installed with a rating of 3000kWh, which means that it is expected to displace an equivalent amount of energy from the grid. The daily expected impact from the system is around 8.2 kWh (3000kWh/365), which implies theoretically that the total amount of energy that might have been consumed in that day without the SHW system would have been roughly 34.7kWh (8.2kWh + 26.5kWh). Of course, the total daily energy consumed varies considerably from day to day, so one must analyze many days from before and after the installation to accurately estimate a representative energy impact from the SHW system.

Similarly, we estimate the power rating for the SHW system to be around 1.34kW (3000kWh/(.25*24*365)). Thus, the maximum hourly average demand reduction from this system is theoretically expected to be around 1.34kW. In this case, the demand for the day without the SHW system would be expected to be 3.74kW (2.4kW+1.34kW). However, like energy, the peak average hourly demand varies considerably from day to day and one must analyze many days before and after installation to accurately estimate a representative demand reduction.

6. VERIFICATION OF INSTALLATIONS AND DETERMINATION OF EXPECTED ENERGY AND DEMAND REDUCTIONS FROM SHW SYSTEMS

We attempted to verify the SHW installations with visual evidence. Since we had the addresses of the homes (but no other personal information), we surveyed them using satellite images that are freely available through Google Earth. Google Earth maintains a library history of all photos for all sites, allowing a user to choose to see the same scene at various times in the past.[2] These pieces of information provided important ground-truth data for the analysis, as we will discuss below.

The photos verified that the installations had occurred. In some cases, a new SHW system was installed where none existed previously. In each of these homes we expected energy and demand reduction to be in ranges that were based on the OG300 ratings. In one case an existing SHW system was removed and replaced with a new one. In this case little change in energy or demand might be expected. In two cases, a two-panel SHW system was replaced with a

[2] The number of historical photos on Google varies from site to site. In recent years photos are have been updated every few months whereas at earlier times, photos were updated every few years. The photo quality and resolution of some of the older photos is considerably inferior to the more recent ones.

single-panel system. In those cases we might expect home energy to rise after installation.

Table 1 presents our estimates of the expected annual energy reduction (in kWh) and demand reduction (in kW) for each of the five homes. We used the OG300 savings rating for each new SHW system as the basis for our expectations.

TABLE 1. EXPECTED ENERGY AND DEMAND REDUCTIONS AFTER SHW INSTALLATION

Home ID	Expected Energy Reduction (kWh/yr)	Expected Power Demand Reduction (kW)
1	2750	1.3
2	-1525	-0.7
3	2750	1.3
4	4097	1.9
5	-1525	-0.7

We had some concern that year-to-year weather variations could produce changes in overall home energy use that exceeded the magnitude of the changes expected from the SHW systems. If sufficiently large, these weather effects could swamp the SHW systems' signatures. However, after studying the heating and cooling degree day statistics for the years under study, we found that the weather variation was much smaller than the size of the expected variation due to the SHW system. Thus, we ignored weather effects.

7. STATISTICAL ANALYSIS TO QUANTIFY THE ENERGY AND DEMAND IMPACT OF SHW SYSTEMS

Our first analytical effort was to simply plot the raw daily energy consumption and power demand over time. We reached two conclusions immediately. First, there appeared to be some visual impact of SHW systems in terms of energy and demand reduction, and these trends seemed to follow roughly the expectations laid out in the table above. Second, there was significant noise in the plots such that the trends were not as clear as we thought they could be after applying some data smoothing techniques.

One of the simplest and most comprehendible smoothing techniques is a moving average. A moving average is the average of a specific number of previous data in a string and using that average to represent a data point at the end of the averaged string. The next data point is computed similarly, and so on.

After experimenting with various moving average periods, a 30-day moving average was selected as the best one for presenting the trends in the energy and demand data for each of the five homes. Shorter averaging periods did not produce sufficient smoothing to clearly demonstrate the trends that we were trying to exemplify and longer averaging periods masked them out.

The plot presented in Fig. 2 shows the 30 day moving averages of energy consumption over the period of record for home 1, the same one represented in Fig. 1. The yellow line shows the approximate date of SHW installation.

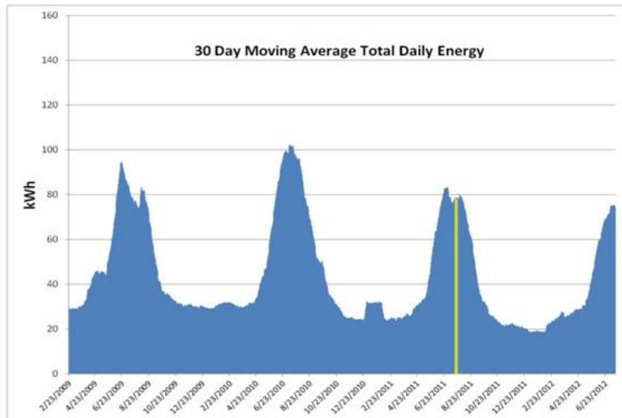

Fig. 2. 30 day moving average of energy consumption vs. time.

As can be seen, there is noticeable reduction in energy use that is evident by comparing both the high marks and low marks for the year after installation to those from before.

Similarly, Fig. 3 shows the 30 day moving averages of power demand for the same home. Again, the reduction in power demand is evident after SHW installation.

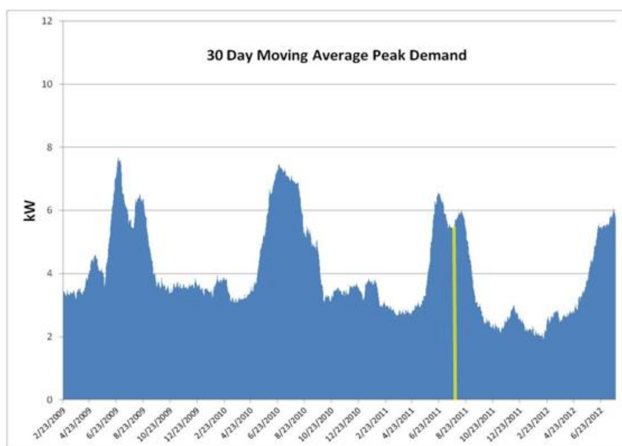

Fig. 3. 30 day moving average of power demand vs. time.

For each home the energy and demand reduction were averaged for the years before and after installation and the differences were taken, thus providing metrics that could be compared to the expected reductions.

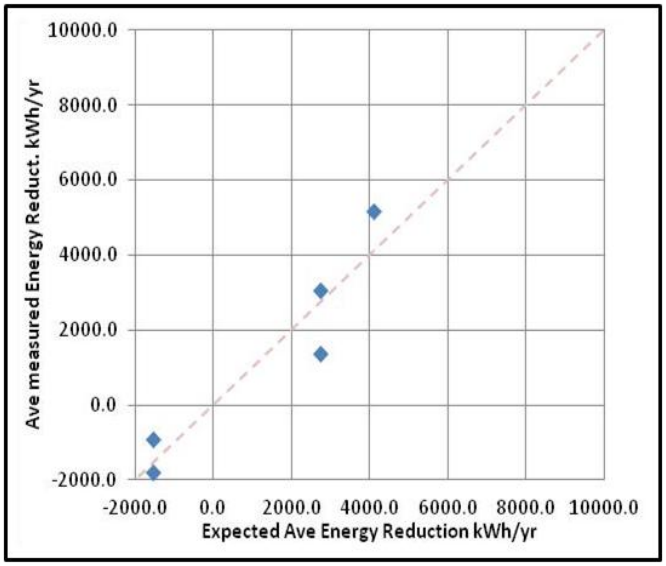

Fig. 4. Measured energy reduction vs. expected energy reduction after SHW installation for five homes under study.

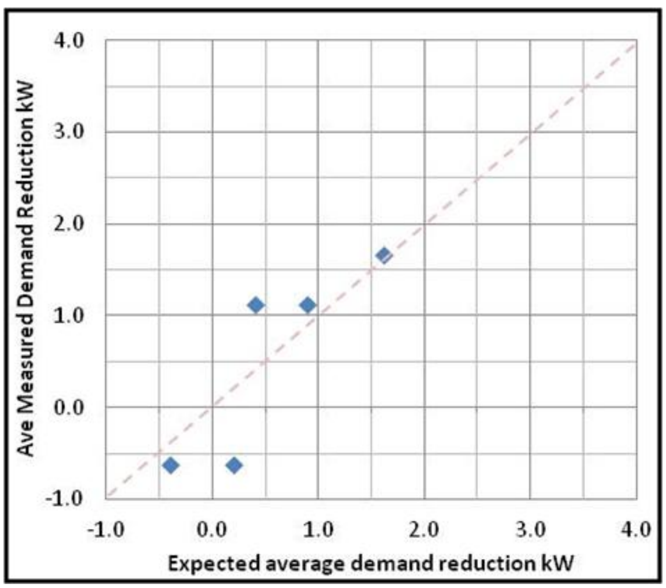

Fig. 5. Measured average power demand reduction vs. expected average power demand reduction after SHW installation for five homes under study.

The graphs in Figs. 4 and 5 show plots of the measured vs. expected energy reduction and measured vs. expected demand reduction for all five homes. The dashed lines show a vector representing perfect relationships.

Considering the variance in the data, the expected and measured values of energy and demand reduction are reasonably well correlated.

8. <u>TEST OF ART'S ABILITY TO DETECT SHW FAILURES</u>

We wished to test ART's ability to recognize failed systems. This capability would allow us to determine the mean time to failure for SHW systems in the sample. However, we could detect no SHW failure in any of our five datasets.

Therefore, to test ART in this capacity we created a simulated failure condition in two of the datasets by copying one year's data from before the installation and appending them to the end of the dataset. For example, a three year period of record with two years prior to SHW installation and one year after became a four year period of record, with the fourth year containing data from one of the first two years.

We then applied ART to the synthetic datasets. An example of one of the applications is shown in Fig. 6 below. The annotations are similar to those in Fig. 1, except that the purple line shows the approximate date of failure for a system. This date was approximately equal to the date at the beginning of the set of appended data. In other words, ART identified the simulated failure with reasonable exactitude.

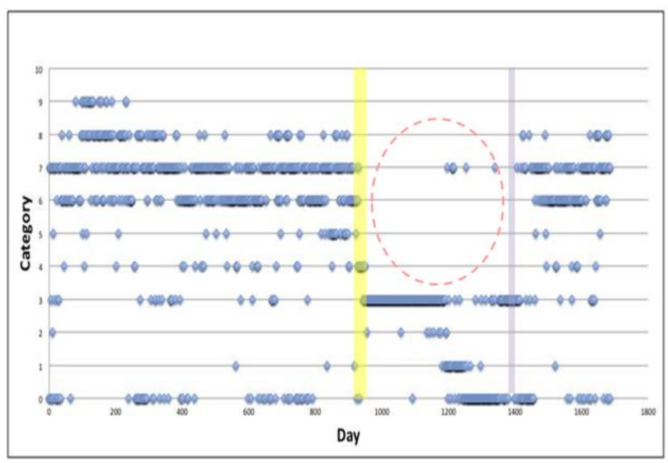

Fig. 6. ART's energy consumption learning categories versus day number for a dataset with a simulated failure.

9. CONCLUSIONS AND RECOMMENDATIONS

Prior to this research it was widely held that SHW energy impact could not be directly measured from whole-house Smart Meter energy records. The results from our analysis show convincingly that by using ART to identify the SHW installation date and subsequently applying simple statistics we can measure the net impact of SHW systems with reasonable clarity. The development of this technology and its application is an important advance in SHW research.

As we noted above, based on our limited sample the measured impact from the SHW systems is about equal to the SRCC energy savings ratings. However, there is a reasonable probability that some of the energy reduction may be due to other factors, such as insulation added to the building shell or upgrades to appliances. SHW installers frequently recommend that energy conservation measures be implemented along with the solar system, and they often perform the upgrades themselves.

However, even if this were the case, it is important information for a utility because the goal of a SHW program is to reduce the load, and these other factors may be a subtle but important positive feature to SHW systems. Importantly, this information has come to light because we analyzed whole-house energy records that incorporate all of the energy saving features associated with the SHW system. Thus, a more complete picture of the SHW systems' impact emerged by studying the total home energy use rather than the SHW systems alone.

While our analysis demonstrates the capability to measure the direct impact of SHW systems on a utility grid, the energy and demand reduction estimates we presented above are based on too few homes to represent the population of fielded SHW systems and should not be used as such. The application of our methodology to a sufficiently large sample of energy records could permit us to provide estimates of SHW impact with a fixed level of precision that would allow their use in a utility's policy, planning and operational activities.

We believe that the ART methodology can be applied to Smart Meter whole-house data to identify SHW failures and to allow the computation of the mean times to failure for these systems, a metric needed by utility resource planners and one that has eluded the SHW community for decades.

We further believe that the analysis of whole-house records is superior to any other methods used to asses SHW grid impact, including computer modeling or field monitoring. Whole-house records contain the most direct measures of SHW activity relative to the grid, which is of prime interest to utilities.

We recommend that additional homes be analyzed with the goal of measuring the energy and demand reduction resulting from SHW systems. Simultaneous with this analysis, we recommend that ART be used to screen for SHW failures. Thus, both studies could be conducted simultaneously using the same data.

Previous work by Menicucci suggests that SHW systems may fail at a constant rate of around 5% annually (4). Assuming two year old systems, which we found in several of our datasets, at least 20 similar homes would have to be examined to expect to find one with a failure.

The desired levels of accuracy in the results can be used to determine the optimal sample size for study. At this time, we estimate that additional analysis should begin with 70 homes. Based on our experience in this pilot study, we expect 20 to be rejected due to data shortcomings of various types. This would leave 50 for the study, which would certainly be sufficient for the failure analysis and would probably be sufficient for determining with reasonable accuracy the mean energy and demand impact of SHW systems.

Fewer homes could be analyzed at proportionally lower cost and less effort, but the results would be proportionally less accurate. We recommend, however, that at least 20 homes be analyzed because that is the minimum number needed to expect to observe one system failure among the samples.

10. REFERENCES

(1) Menicucci, D., Report on the Analysis of Field Data Relating to the Reliability of Solar Hot Water Systems, Sandia National Labs report, SAND2011-4759, July, 2011.

(2) Carpenter, G. A., Grossberg, S., and Rosen, D. B., "Fuzzy ART: Fast stable learning and categorization of analog patterns by an adaptive resonance system," Neural Networks, vol. 4, pp. 759–771, 1991.

(3) Menicucci, D., "Solar Thermal Power Rating," ReFocus Journal, March, 2006.

(4) Menicucci, D., Assembly and Comparison of Available Solar Hot Water System Reliability Databases and Information, Sandia National Labs report, SAND2009-2757, May 2009.