

# Algorithms for Green Buildings: Learning-Based Techniques for Energy Prediction and Fault Diagnosis

*Daniel Holcomb  
Wenchao Li  
Sanjit A. Seshia*



Electrical Engineering and Computer Sciences  
University of California at Berkeley

Technical Report No. UCB/EECS-2009-138

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-138.html>

October 15, 2009

Copyright © 2009, by the author(s).  
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

# Algorithms for Green Buildings: Learning-Based Techniques for Energy Prediction and Fault Diagnosis

Daniel Holcomb      Wenchao Li      Sanjit A. Seshia  
EECS Department  
UC Berkeley  
{holcomb,wenchaol,sseshia}@eecs.berkeley.edu

## Abstract

*We consider two problems in the design and operation of energy-efficient buildings. The first is the prediction of energy consumption of a building from that of similar buildings in its geographical neighborhood. The second problem concerns the localization of faults in building sub-systems with a focus on faults that lead to anomalous energy consumption. For both problems, we propose algorithmic techniques based on machine learning to address them. Simulation results using EnergyPlus show the promise of the proposed methods.*

## 1 Introduction

Computation is being widely embedded within our environment, such as systems in health care, transportation, agriculture, and buildings. There are several societal challenges that can be addressed with embedded distributed computing, of which perhaps the defining challenge of our time is to improve the efficiency of energy usage. In particular, the problem of improving the energy-efficiency of buildings is particularly important.

A recent report by the American Physical Society (APS) [13] describes the need for energy-efficient buildings:

“Americans spend 90 percent of their time indoors... In 2006, buildings, more than 118 million residential and commercial structures, were responsible for 39 percent of the nation’s primary energy consumption... Yet a large fraction of the energy delivered to buildings is wasted because of inefficient building technologies.”

The report estimates that it is possible to achieve as much as a 70% savings in energy consumption in new buildings by 2030. Perhaps more importantly, there is a need to *retrofit existing buildings* to improve their energy efficiency.

The work we report on in this paper is motivated by two findings in the report, paraphrased below:

- There is a need for *integrated design*, a process in which all design variables are considered together to arrive at the optimal design which meets user requirements and minimizes energy consumption.
- Advanced technologies are needed to support maintenance of building automation systems, including techniques for *diagnostics, fault detection and control in real time*.

Consider the first finding. In order to optimize energy consumption, it is first necessary to be able to predict how the energy consumption varies as a function of design variables and environmental parameters. The challenge is that many of these variables are unknown, and measuring them can be quite difficult or even impossible. We

address this point in Section 2 of this report by presenting *a new approach to predict building energy consumption based on the consumption of other similar buildings in the neighborhood*.

The second finding highlights the need for fault detection and diagnostics. A second contribution of this report is an investigation of *learning-based methods to localize faults that lead to anomalous energy consumption*.

Our experimental evaluation relies on the use of the EnergyPlus [2] simulator for building energy consumption. We include a discussion of related work on energy prediction and fault diagnosis within the respective sections.

## 2 Building Energy Prediction

A number of approaches have been proposed to predict the energy consumption of a building. They can be largely grouped into two main categories - those based on time series and those that are equation-based. Time-series prediction relies on looking at the history of energy consumption of a building and corresponding weather data, without considering the building parameters, and extrapolating that data into the future [8]. However, this method cannot be applied to estimating energy consumption of a new building since there would be no historical data to learn from. On the other hand, equation-based tools begin with a physics model such as the building structure and heat transfer equations and simulate the model using finite element method [2]. However, there are two main drawbacks of this approach. First, equation-based techniques can be computationally expensive and thus may not scale to a large system. Second, since the method heavily depends on the accuracy of the model, it suffers from environmental uncertainties such as weather conditions.

We propose a different approach for predicting energy consumption of a residential building that takes advantage of the availability of structurally similar buildings in its neighborhood. This approach has the advantage of taking environmental uncertainties out of the equation since buildings in the same neighborhood would face very similar weather conditions. It is often true that buildings in the same neighborhood are constructed at the same time, and thus are quite similar in terms of their layout and design parameters. So the idea is to pool such parameters from those buildings and construct a model that can be used to estimate the energy consumption of the target building. We formally define our approach in the rest of the section.

For a building  $B$ , we have a vector of known parameters  $\vec{x}$ . These can be static parameters such as the thickness of a wall. They can also be statistical estimates such as population density. In addition, we can obtain control parameters such as minimum temperature allowed by the HVAC system. Our training set  $X$  is a  $n \times p$  matrix where row  $i$  is the  $\vec{x}_i$  for building  $B_i$ . The corresponding response variables are a column vector  $Y$  that contains the building  $B_i$ 's daily total energy consumption  $E_i$ . The objective is then to learn a model from  $X$  and  $Y$  and given a new set of parameters  $\vec{x}_n$ , we want to estimate the corresponding daily energy consumption  $E_n$ .

### 2.1 Energy Estimation

The building model that we use for energy estimation is shown in Figure 1. The parameters that we consider in this model are locations of internal walls, ceiling heights, min/max temperature allowed by control, heating/cooling air temperature, lighting, density of people, fiberglass insulation thickness for exterior walls, roof fiberglass insulation thickness, window thickness and roof solar absorbance. We do not consider loads from electrical appliances since this can be treated as adding a constant to the energy consumption.

We evaluate our approach with three statistical learning techniques: artificial neural network (ANN), multilinear regression and support vector regression (SVR). We also perform our experiments on four different days: January 1, April 1, July 1 and October 1 respectively. These days are representative of the four different seasons in the United States.




Figure 1: A three room building model.

### 2.1.1 Artificial Neural Network

Artificial neural networks have been shown to be flexible in mapping nonlinear models from inputs to outputs [5]. However, while it places few limitations on the input-output relations, it requires much data to reliably find the correct relation. It is also notorious for overfitting (fitting noise). Readers can find a more detailed study of the application of ANN to buildings in Dodier’s thesis [3]. We use a non-recurrent (acyclic) network in our analysis. Mathematically, the input-output relation for our network is the following.

$$f(x, w) = g_{output}(b_o + \sum_{k=1}^K v_k g_{hidden}(b_h + \sum_{l=1}^L u_{kl} x_l)) \quad (1)$$

where  $g_{hidden}$  is the activation function for the hidden nodes,  $g_{output}$  is the output function for the output nodes,  $b_o$  and  $b_h$  are biases, and  $u_{kl}$  and  $v_k$  are weights.

Figure 2 shows the network that we use for our experiments with one hidden layer with  $K = 20$  nodes and one output layer with a single node. There are  $L = 11$  inputs. The activation functions that we use are hyperbolic tangent for hidden nodes and a linear transfer function for the output nodes. We train the network using the feed-forward back-propagation method [4].




Figure 2: Artificial Neural Network

### 2.1.2 Multilinear Regression

The multilinear regression model with  $p$  independent variables is the following:

$$y_i = \beta_0 + \beta_1 x_{1i} + \beta_2 x_{2i} + \dots + \beta_p x_{pi} + \epsilon_i \quad (2)$$

where  $\epsilon_i$  are independent mean zero error terms. The coefficients  $\hat{\beta}$  of the best fit in terms of least square can be found by

$$\hat{\beta} = (X^T X)^{-1} X^T Y \quad (3)$$

### 2.1.3 Support Vector Regression

We use the  $\epsilon$ SVR method devised by Vapnik et al. [14]. Given training data, the idea is to find a function  $f(x)$  that has at most  $\epsilon$  deviation from the actually obtained targets  $y_i$  for all the training data and at the same time as flat as possible. The corresponding optimization is the following

$$\begin{aligned} & \text{minimize } \frac{1}{2} \|w\| + C \sum_{i=1}^n (\xi_i + \xi_i^*) & (4) \\ & \text{subject to } \begin{cases} y_i + \langle w, x_i \rangle - b \leq \epsilon + \xi_i \\ \langle w, x_i \rangle + b - y_i \leq \epsilon + \xi_i^* \\ \xi_i, \xi_i^* \geq 0 \end{cases} \end{aligned}$$

Flatness here means we would like to get a small  $w$ . The constant  $C > 0$  determines the trade-off between the flatness of  $f$  and the amount up to which deviations larger than  $\epsilon$  are tolerated.  $\xi_i$  and  $\xi_i^*$  are slack variables for the optimization problem. In our experiments, we use a linear kernel and  $C = 10000$ .

### 2.1.4 Experimental Results

We first perform some initial processing on  $\vec{x}$  including taking the sine of the building’s orientation and normalizing the other columns. We perform experiments on 3 sets of buildings, each with 50 as the training set and 50 as the test set. The data are generated by EnergyPlus [2]. The building model that we use is the 3-room building in Figure 1. The weather for the building is historical data from Chicago, Illinois.




Figure 3: Accuracy with ANN

The first set contains structurally identical buildings. The second set contains structurally similar buildings (small variations in internal and external walls and changes in orientation). The last set contains structurally different buildings (large variations of the aforementioned parameters). Figures 3, 4, 5 show the prediction accuracy on Jan 1 for using ANN, multilinear regression and SVR. In terms of mean squared error, SVR beats multilinear regression by a small margin but both significantly outperforms ANN. This is because ANN requires a lot of data to correctly model the input-output relations of a system. On the other hand, while SVR is slightly more accurate than multilinear regression, it can be more computationally expensive.




Figure 4: Accuracy with Multilinear Regression




Figure 5: Accuracy with SVR

Table 1, 2, 3 show the comparisons for the three methods for the 4 different days in terms of mean squared error (in  $10^6$  Joules).

The results show that our approach is in fact robust to small changes in building structures. We can accurately estimate the new building's energy consumption by using only 50 similar buildings and pooling 11 building parameters. However, if the buildings are significantly different, we are not able to estimate the energy consumption reliably.

As a future direction, it would be interesting to use an energy optimization technique, such as GenOpt [15], in conjunction with energy prediction so as to improve the energy efficiency of similar buildings (or apartments) in the same neighborhood.

Same	ANN	Multilinear Regression	SVR
Jan. 1	13289	1669	1347
Apr. 1	9421	918	765
Jul. 1	2302	92	83
Oct. 1	3173	291	280

Table 1: Structurally Identical Buildings

Similar	ANN	Multilinear Regression	SVR
Jan, 1	24613	4755	7404
Apr, 1	15340	1797	1622
Jul, 1	4174	180	241
Oct, 1	8224	542	859

Table 2: Structurally Similar Buildings

### 3 Fault Detection and Diagnosis

Faults in building energy systems are often a cause of reduced energy efficiency. The goal of fault detection and diagnosis (FDD) is (1) detecting that a fault occurs, and (2) generating diagnostic information to localize and identify the fault. FDD can be either model-based or history-based, and can be applied at the component scale up to the whole building scale.

Our FDD approach is demonstrated on the building shown in Figure 1 in Section 2; it is an L-shaped single floor rectangular building with a 40 ft south wall, 40 ft west wall, and 10 foot high ceilings. The building has a single window, with walls made of stucco over brick, and a 1/2 inch stone roof. This building is derived from one of the example buildings included with the EnergyPlus software. The building has a variable-air-volume (VAV) reheat system, and three distinct thermal zones. The cooling for all zones is provided by a central chilled water coil, and zones are heated individually. The weather for the building is historical data from Chicago, Illinois. Temperature is controlled according to dual setpoints with a 4C deadzone between.

#### 3.1 Related Work

Piette et al. [11] demonstrate a system level Information Monitoring Diagnostic System (IMDS) that logs 57 sensors including temperature, power, pressure, and flow sensors. The IMDS does not perform FDD, but includes data visualization software to facilitate manual FDD.

Choi et al. [1] demonstrate model-based FDD in an HVAC chiller component using a three step process of fault detection, fault isolation, and fault severity estimation. 28 sensors are used to monitor input and output temperatures, pressures, flow rates, power, and enthalpy. 5 different classes of faults are applied. Gaussian measurement noise is added to measurements. and a low-pass (0-3Hz) filter is used to reject noise. Assuming that the distribution of unfaulted sensor readings is known, a generalized likelihood ratio (GLR) test is used to detect the occurrence of a fault. Once a fault is detected, multiway dynamic PCA (MPCA), multiway partial least squared (MPLS) regression, and SVMs are used for fault classification.

Rossi et al. present an FDD scheme for rooftop air conditioner units. The interesting aspect of this work is that fault diagnosis is only performed when the unit is in steady state; each time a measurement is taken, a classifier is used to determine whether or not the system is in steady state. When the system is in steady state, residuals are used for error-detection [12].

Liang and Du demonstrate a model-based multi-layer FDD in HVAC systems [9]. Faults are detected by ana-



Different	ANN	Multilinear Regression	SVR
Jan, 1	660200	236910	190297
Apr, 1	296190	111690	133310
Jul, 1	45653	19355	25214
Oct, 1	102060	42947	47335

Table 3: Structurally Different Buildings

lyzing residuals; SVM classifiers are used for diagnosis <sup>1</sup>. A separate SVM is trained for each of 4 hypotheses (3 different faulted states, and the unfaulted state); The features used include supply air and water temperatures and control signals. Given that this approach assumes availability of replay-generated fault residuals, it is unclear how successful it would be in a real building, where replay is not possible.

Katipamula et al. detect and diagnose outdoor air economizer faults using a decision tree [6][7]. The decision tree is generated manually based on “rules derived from engineering models of proper and improper air-handler performance”. The decision variables are the building energy, air temperatures, air humidities, fan schedules, and damper position.

Like several related works, our approach is a two-step process of error detection and fault diagnosis. However, our work is unique among automated diagnosis techniques in using whole building energy consumption as the error detection criterion.

### 3.2 Problem Formulation

Our approach uses energy prediction for error detection, and uses a decision tree for fault diagnosis. It assumes the existence of an accurate building model that can be used for fault injection experiments. A building  $B$  is viewed as a mapping from environment and fault-state to a set of sensor readings.  $B : S_E \times F \rightarrow S_B$ .

- $F \in \{f_1, \dots, f_{10}\}$  is a random variable for the fault state of the building, as described in Table 4.
- $S_E \in \mathcal{R}^M$  is a variable representing measurable environmental quantities;  $S_E$  is independent of  $F$ .
- $S_B \in \mathcal{R}^N$  is a variable representing the measurable state of building sensors.

Observations of the building are made hourly. Each observation of the building provides  $S_E$  and  $S_B$ .

F	description
$f_1$	VAV1 damper cannot open fully
$f_2$	VAV2 damper cannot open fully
$f_3$	VAV3 damper cannot open fully
$f_4$	VAV1 damper cannot close
$f_5$	VAV2 damper cannot close
$f_6$	VAV3 damper cannot close
$f_7$	VAV fan efficiency decreased
$f_8$	HW circ pump max-flow rate drop
$f_9$	CW circ pump max flow rate drop
$f_{10}$	cond. circ pump max flow rate decreased

Table 4: Description of injected fault types.

<sup>1</sup>the residuals are generated by replaying a model in faulted and unfaulted states.

### 3.2.1 Error Detection

We define an error to occur if the the hourly observed energy consumption ( $E_{OBS}$ ) does not closely match the predicted energy consumption ( $E_{PRED}$ ). More precisely, the error criterion is given by Equation 5, where  $\epsilon$  is the hourly energy threshold for error detection.

$$error = |E_{OBS} - E_{PRED}| > \epsilon \tag{5}$$

Error detection is thus performed by comparing  $E_{OBS}$  against  $E_{PRED}$ . Depending on the environment, an injected fault may or may not cause an error.

In this work,  $E_{PRED}$  is obtained from *fault-free simulation*; in future work, it could be obtained from a regression model. This model would be trained using historical environment data ( $S_E$ ) and energy consumption ( $E_{OBS}$ ) from times when the building was known to be fault-free.

The histogram of Figure 6 shows the distribution of  $E_{OBS} - E_{PRED}$  for each fault type. To generate this data, each of the 10 different fault-types is injected on 200 randomly chosen days, yielding 4800 hourly observations for each fault. Each time a fault is injected, its severity is chosen uniformly from a range. Because historical weather data is being used, the random dates cause the  $S_E$  to vary.

A fault can either increase or decrease the energy consumption of a building. For example, a damper that cannot close will often cause an increase in energy consumption, while a damper that cannot open will reduce the energy consumption. In cases where the energy decreases due to a fault, it is likely that comfort or air quality considerations would suffer, so the fault must still be diagnosed.

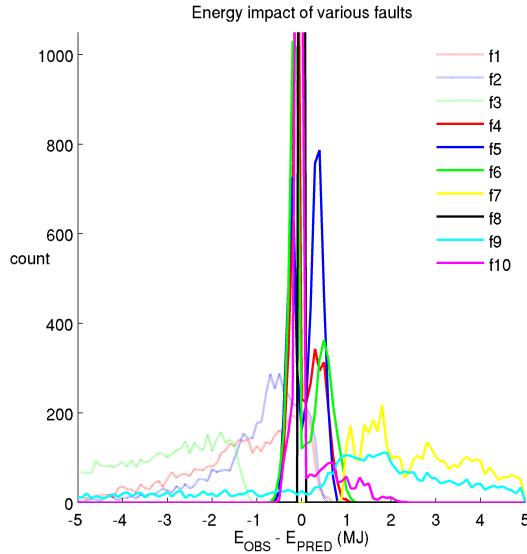


Figure 6: Histogram of  $E_{OBS} - E_{PRED}$  for each fault type.

In this, work we nominally set  $\epsilon$  to 360KJ. Table 5 shows how number of detected errors and diagnosis accuracy for other values of  $\epsilon$ . At the 360KJ threshold, fault  $f_8$  can never cause a detectable error (see Fig. 6), so it does not factor into diagnosis. Among the faults to be diagnosed,  $f_{10}$  causes the fewest errors, exceeding the threshold in just 791 of the 4800 fault injections.

$\epsilon$ (KJ)	accuracy (%)	total errors
100	52.13	35382
200	76.26	31246
300	81.79	28143
360	76.86	26584
400	78.59	25536
500	80.41	23364

Table 5: Accuracy and total errors for various  $\epsilon$ .

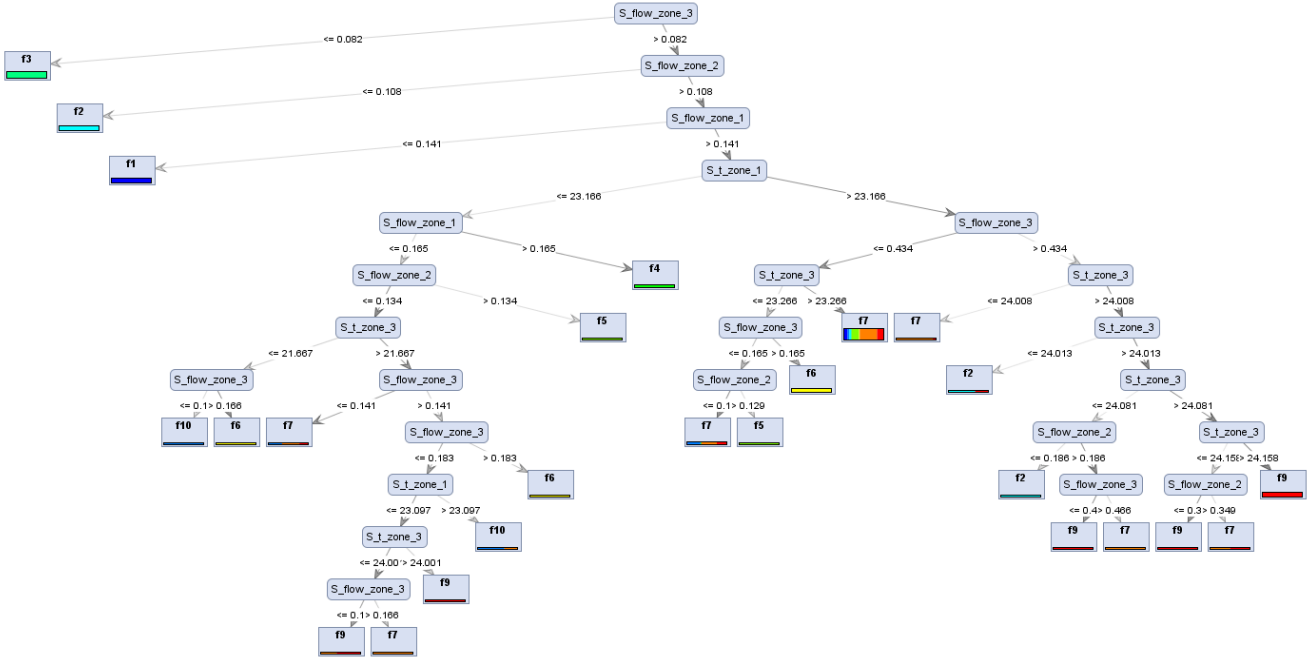


Figure 7: Diagnostic decision tree created using  $S_E$  and  $S_B$ .

### 3.2.2 Fault Diagnosis

If an error is detected in an hourly observation, we need to diagnose the fault type. In this work, we learn a decision tree model and use it to diagnose the fault type. The inputs to the decision tree are environment  $S_E$  and the sensor readings  $S_B$ ; the leaves of the decision tree are the fault-types  $F$ .

The decision tree is automatically trained from labeled observations, using the software package RapidMiner [10]. The training and test data are generated by simulating the faults of Table 4. In generating the data, it is assumed that the faults naturally occur uniformly at random. Errors are not uniform, as some faults are more likely to cause errors than others. Once the fault type is chosen, the fault severity is also chosen randomly. Only observations that contain errors (per Eq. 5) are used as training and test cases. Using half of the data for training the decision tree and half for testing, the accuracy of diagnosis is 75.8% (Tab. 6.).

### 3.3 Efficient Sensor Selection for Diagnosis

For a building with some specified fault universe, one might wish to select an efficient subset of sensors to install in order to diagnose faults. As a by-product of creating the decision tree, the information gain of each sensor is quantified. Installing only the sensors with the highest information gain will lead to most efficient fault

True:	f1	f10	f2	f3	f4	f5	f6	f7	f9
f1:	1259	0	0	0	0	0	0	0	0
f10:	0	2	0	0	0	0	0	5	2
f2:	0	0	1414	0	0	0	0	3	2
f3:	0	0	0	2423	0	0	0	0	0
f4:	0	0	0	0	522	0	0	0	0
f5:	0	0	0	0	0	77	0	0	0
f6:	0	0	0	0	0	0	760	0	0
f7:	466	383	256	0	48	880	89	2405	961
f9:	0	3	0	0	0	0	0	4	1327

Table 6: Diagnosis results. The column indicates the true fault that was injected, and the row indicates the diagnosed fault.

sensor	information gain
damper flow, zone 3	1.0
damper flow, zone 2	0.8806
damper flow, zone 1	0.8229
temp, zone 1	0.3788
temp, zone 3	0.3573
temp, zone 2	0.3409
windspeed, outdoor	0.0022
occupants, zone 1	0.0021
occupants, zone 2	0.0021
occupants, zone 3	0.0021
time of day	0.0011
temp, outdoor	0.0008
time of year	0.0

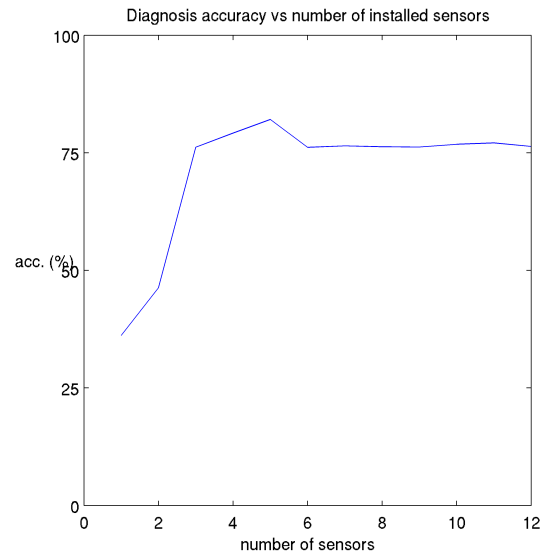


Figure 8: Diagnosis accuracy as a function of number of sensors, with sensors ordered by decreasing information gain.

diagnosis.

The information gain of each sensor variable is shown at left in Figure 8. Starting with no sensors, we add sensors in order of decreasing information gain, and evaluate diagnosis accuracy after each additional sensor. Many of the sensors contribute minimally to diagnosis, as using just 3 sensors gives comparable results to using all 12. In this case, it turns out that the top three sensors are all flow sensors, and this is likely an artifact of the chosen fault universe consisting largely of stuck damper faults. If a different fault universe were of interest, then a different set of sensors might be selected.

### 3.4 Conclusion

This work explores whether anomalous energy consumption and basic sensors might be used to guide whole-building error detection and fault diagnosis. With a set of sensors that is well-suited to the faults of interest, diagnosis may be possible. A consideration when using energy consumption for error detection is that many faults have little or no impact on energy consumption; if these faults have other undesirable effects, then an alternative approach for detection is needed.

## Acknowledgments

The work described in this report has benefited from valuable discussions on the topic with several people and from the EECS 290-N course taught at Berkeley in Spring 2009. In particular, we thank Jaijeet Roychowdhury, Alberto Sangiovanni-Vincentelli, and Michael Wetter. This work was supported in part by the Gigascale Systems Research Center (GSRC), a Hellman Family Faculty Fund award, and an Alfred P. Sloan Research Fellowship.

## References

- [1] K. Choi, SM Namburu, MS Azam, J. Luo, KR Pattipati, and A. Patterson-Hine. Fault diagnosis in HVAC chillers. *IEEE Instrumentation & Measurement Magazine*, 8(3):24–32, 2005.
- [2] Drury B. Crawley, Linda K. Lawrie, Frederick C. Winkelmann, W. F. Buhl, Y. Joe Huang, Curtis O. Pedersen, Richard K. Strand, Richard J. Liesen, Daniel E. Fisher, Michael J. Witte, and Jason Glazer. EnergyPlus: A new-generation building energy simulation program. *Energy & Buildings*, 33:319–331, April 2001.
- [3] R.H. Dodier. Statistical properties of neural networks, with application to building energy prediction. Master’s thesis, University of Colorado, 1995.
- [4] M. T. Hagan and M. B. Menhaj. Training feedforward networks with the Marquardt algorithm. *IEEE transactions on Neural Networks*, 5(6):989–993, 1994.
- [5] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- [6] S. Katipamula, M.R. Brambley, N.N. Bauman, and R.G. Pratt. Enhancing building operations through automated diagnostics: Field test results. In *Proceedings of the Third International Conference for Enhanced Building Operations*, 2003.
- [7] S. Katipamula, R.G. Pratt, D.P. Chassin, Z.T. Taylor, K. Gowri, and MR Brambley. Automated fault detection and diagnostics for outdoor-air ventilation systems and economizers: Methodology and results from field testing. *Transactions-American Society Of Heating Refrigerating And Air Conditioning Engineers*, 105:555–567, 1999.
- [8] J.F. Kreider and J.S. Haberl. Predicting hourly building energy use: the great energy predictor shootout-overview and discussion of results. *ASHRAE Transactions*, 100(2):1104–1118, 1994.
- [9] J. Liang and R. Du. Model-based fault detection and diagnosis of HVAC systems using support vector machine method. *International Journal of Refrigeration*, 30(6):1104–1114, 2007.
- [10] Ingo Mierswa, Michael Wurst, Ralf Klinkenberg, Martin Scholz, and Timm Euler. Yale: Rapid prototyping for complex data mining tasks. In Lyle Ungar, Mark Craven, Dimitrios Gunopulos, and Tina Eliassi-Rad, editors, *KDD ’06: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 935–940, New York, NY, USA, August 2006. ACM.
- [11] M.A. Piette, S.K. Kinney, and P. Haves. Analysis of an information monitoring and diagnostic system to improve building operations. *Energy & Buildings*, 33(8):783–791, 2001.
- [12] T.M. Rossi and J.E. Braun. A statistical, rule-based fault detection and diagnostic method for vapor compression air conditioners. *International Journal Of Heating Ventilation Air Conditioning And Refrigerating Research*, 3:19–37, 1997.

- [13] American Physical Society. Energy future: Think efficiency. <http://www.aps.org/energyefficiencyreport/>, September 2008.
- [14] V. Vapnik, S.E. Golowich, and A. Smola. Support vector method for function approximation, regression estimation, and signal processing. *Advances in Neural Information Processing Systems 9*, 1997.
- [15] M. Wetter. GenOpt - a generic optimization program. In *Proc. of the Seventh International IBPSA Conference*, pages 601–608, 2002.