




Extracting Operating Modes from Building Electrical Load Data

Preprint

Stephen Frank
Colorado School of Mines

Luigi Gentile Polese, Emily Rader,
Michael Sheppy, and Jeff Smith
National Renewable Energy Laboratory

Presented at the 2011 IEEE Green Tech Conference
Baton Rouge, Louisiana
April 14 – 15, 2011

Conference Paper
NREL/CP-5500-49636
January 2012

Contract No. DE-AC36-08GO28308

Available electronically at http://www.osti.gov/bridge

Available for a processing fee to U.S. Department of Energy
and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information

P.O. Box 62
Oak Ridge, TN 37831-0062
phone: 865.576.8401
fax: 865.576.5728
email: mailto:reports@adonis.osti.gov

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
phone: 800.553.6847
fax: 703.605.6900
email: orders@ntis.fedworld.gov
online ordering: http://www.ntis.gov/help/ordermethods.aspx

Cover Photos: (left to right) PIX 16416, PIX 17423, PIX 16560, PIX 17613, PIX 17436, PIX 17721

Printed on paper containing at least 50% wastepaper, including 10% post consumer waste.

Extracting Operating Modes from Building Electrical Load Data

Stephen Frank
Division of Engineering
Colorado School of Mines
Golden, Colorado, United States of America
stfrank@mymail.mines.edu

Luigi Gentile Polese, Emily Rader,
Michael Sheppy, and Jeff Smith
Electricity, Resources, & Building Systems Integration
National Renewable Energy Laboratory
Golden, Colorado, United States of America
luigi.gentile.polese@nrel.gov

***Abstract*—Empirical techniques for characterizing electrical energy use now play a key role in reducing electricity consumption, particularly miscellaneous electrical loads, in buildings. Identifying device operating modes (mode extraction) creates a better understanding of both device and system behaviors. Using clustering to extract operating modes from electrical load data can provide valuable insights into device behavior and identify opportunities for energy savings. We present a fast and effective heuristic clustering method to identify and extract operating modes in electrical load data.**

Keywords–buildings, clustering, miscellaneous electrical loads, mode extraction, system characterization.

I. INTRODUCTION

Miscellaneous Electrical Loads (MELs) are the fastest growing load category in buildings and account for the largest portion of commercial and residential electricity consumption. MELs are defined as all building electrical loads apart from heating, ventilation, and air-conditioning (HVAC), space lighting, and water heating [1]. The U.S. Department of Energy (DOE) estimates that MELs will grow from 31% of total commercial building primary energy use in 2006 (1610 TWh or 5.5 Quads) to 43% in 2030 (3140 TWh or 10.7 Quads); residential buildings have similar statistics. Since commercial and residential buildings consume 40% of U.S. primary energy (29300 TWH or 100 Quads in 2006), MELs account for around 10% of all energy used in the United States [2], [3].

MELs represent a growing opportunity for energy savings. Characterizing MEL behavior is a key step toward energy reduction. Mode extraction—the identification and quantification of operating states (modes)—is an important technique for load characterization. Data collection and mode extraction for MELs are difficult to implement, however, because many end use subcategories fall under the umbrella of "miscellaneous" load. Several studies have specifically addressed MELs, particularly plug loads [1], [4]-[7], but few sources of load data are available for specific MELs.

Targeted metering of individual MELs can alleviate the lack of available load data. As part of a larger DOE effort, the National Renewable Energy Laboratory (NREL) is conducting a study of commercial MELs. The study targets a large retail store with a wide variety of space types, including general merchandise, grocery, restaurant, pharmaceutical, and medical services. NREL has catalogued more than 750 distinct MEL devices and metered approximately 240 of these, recording power, voltage, and energy use. Characterizing and modeling MEL behavior are primary outcomes of the study, both to identify energy saving opportunities and to improve energy simulations. A rapid, reliable method of identifying and extracting MELs power consumption characteristics is needed to create energy models. Here, we present a fast and effective heuristic clustering technique for extracting operating modes from electrical load data.

II. RELATED WORK

Several studies have explored MEL operation, including extraction of operating modes. A recent study of MELs in U.S. commercial buildings [1] used building-level electrical data to estimate national MELs energy use by device and building type. This study is among the most rigorous in combining independent MELs data sets, but its results have significant uncertainty because many analysis inputs are assumptions, averages, or best estimates. A large-scale residential study in Europe [4] employed an integrated approach consisting of monitoring studies, household surveys, and detailed audits. The study focused on characterizing and reducing standby power consumption as defined in the IEC62301 standard [8], however, as opposed to overall power consumption.

More limited studies have targeted specific MEL categories via field monitoring. One study examined a subset of office MELs via intensive monitoring over a two week period [5]. A similar study that focused on residential uses in Germany [6] analyzed and characterized power modes and operating times of information and communication devices according to four pre-defined operating modes: normal, standby, off-mode, and off. Another study employed non-intrusive methodology consisting of an inventory and a night-time audit of the power status of office equipment and devices [7]; it recorded data on the types, power states and power save delay settings of office equipment in a commercial building environment. This study consisted only of a snapshot at a single point in time, however, and therefore yielded no information about the time spent in each power state.

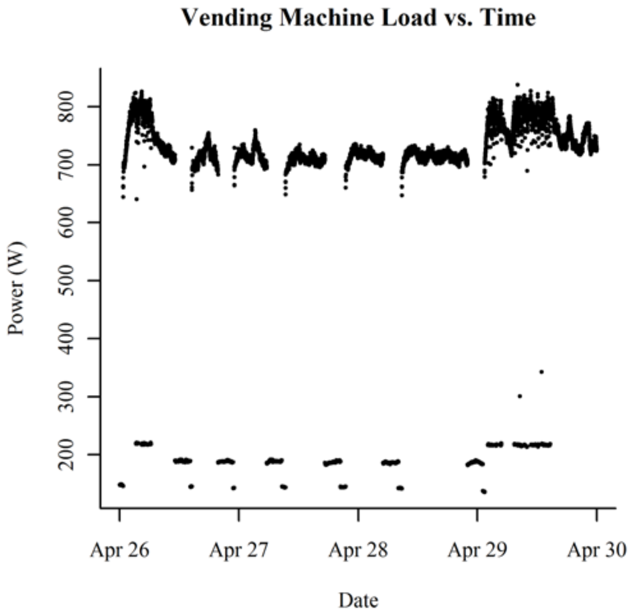

Figure 1. Example of time series load data recorded from a vending machine used as an input for mode extraction via clustering.

Non-intrusive load monitoring (NILM) offers a different approach to load characterization, for both MELs and traditional loads. NILM uses a single, centralized instrument to monitor aggregated electricity consumption. Computer software then disaggregates individual loads from the overall signal, recording device behavior in the process. The simplest and least expensive NILM techniques measure changes in real and reactive power levels and require only low frequency sampling [9]-[11]. More complex techniques rely on harmonic analysis to distinguish loads; these require costlier hardware and more sophisticated software [12], [13]. NILM can reliably distinguish on/off loads, but loads with multiple or variable states present a much greater challenge [9], [14], [15]. NILM is less disruptive and less costly to deploy than a plug-level metering and can track mobile loads precisely. However, NILM must be preceded by analysis of the individual loads to be monitored because NILM algorithms can disaggregate individual loads only by comparison with known load behavior. Therefore, NILM is not a suitable first step in load characterization.

III. Mode Extraction via Clustering

Cluster analysis, or clustering, is the process of dividing a data set into logical groupings (clusters) based on similarity criteria. Mode extraction for electrical loads may be viewed as a clustering problem: each power mode represents a distinct cluster within the load data that must be identified and extracted. The resulting modes (clusters) may then be further analyzed to determine characteristics such as duty cycle, average mode power, and expected mode transitions.

For example, the vending machine power data of Figure 1 suggest four operating modes: a dominant mode near 700 W and three distinct but closely spaced modes in the range of 150-250 W. Clustering the data provides a method for detecting and quantifying these power modes.

A. Data Set Challenges

Electrical load data present several challenges that make effective clustering difficult. Long-term time series of electrical loads may contain several hundred thousand data points, recording steady-state device operation, transient conditions, and noise. Many assumptions that give tractability to other clustering problems do not hold for electrical load data: the number of clusters is not known a priori, they may not conform to Gaussian distributions, and they may be of different densities. Electrical loads, particularly MELs, differ widely with respect to key characteristics, including magnitude, characteristic load profile, and cycling frequency. This load diversity creates difficulty during clustering algorithm initialization, as little information is available in advance. For loads encountered in the field, even basic nameplate data are often unavailable or unclear.

Instrumentation imposes further challenges because a cost versus data quality tradeoff is inherent in widely deployed sensor networks. Data transmission and storage capabilities limit time resolution and the number of parameters recorded. Measured data are discretized to the precision of the instrument. Poor-quality instrumentation can also create excess noise, which must be discarded during mode extraction.

B. Algorithm Requirements

In characterization studies involving many loads, clustering must be performed on a large number of data sets representing diverse devices. The algorithm employed should operate with minimal user input, as manual initialization is tedious when many data sets must be processed. The algorithm should also be robust with respect to the challenges identified in section III-A. An ideal clustering algorithm requires no initialization, is fast, is computationally efficient for large data sets, identifies and removes noise in the data, identifies clusters of arbitrary shapes, and accurately determines cluster boundaries.

No presently available clustering algorithm fulfills all these requirements. In practice, the user must at minimum set global defaults for the algorithm and specify which measurements to use. The case where default parameters may be automatically computed from data set characteristics represents a suitable level of automation. Also, clustering becomes more tractable when it is performed using only a few dimensions. Here, we limit cluster analysis to device power only.

C. Clustering Algorithms

Numerous cluster analysis techniques are available, many of which are summarized in [16] and [17]. Only a few common partitioning techniques are discussed here, as many clustering techniques are too computationally expensive for application to large data sets [18]. Centroid-based techniques, such as K-means, are a subclass of partitioning techniques that classify the data into k clusters based on the distance to k central points, where k is an input variable. Such techniques must be iterated for differing values of k when the number of clusters is not known. This greatly increases algorithm run time.

Another class of partitioning methods applies the expectation-maximization (EM) algorithm for statistical modeling to finite mixtures, for instance, a set of Gaussian distributions [19]. EM-based models must select both the number of clusters and an appropriate statistical model a priori. If this information is not known, the technique may perform very poorly. EM-based algorithms are more computationally expensive than centroid-based techniques, as they require complex statistical calculations.

Density estimation partitioning techniques define clusters as regions of high density within the data. In contrast to the previous techniques, density estimation requires no a priori information about the number or shape of data clusters. The DBSCAN algorithm and its variants use a nearest-neighbor definition of density to perform clustering [18], [20], [21]. Except in cases where Fast Nearest Neighbor algorithms are available, the computational cost of these techniques scales with the square of the data set size [18].

IV. The Histogram Heuristic Approach

Because density estimation requires little input prior to clustering, it readily lends itself to automation; however, density-based approaches are computationally expensive. We developed a rapid approach called the histogram heuristic (HH) technique, which uses a histogram to estimate data density. The technique is based on the concept of visual classification of electrical load data and automates an otherwise tedious manual process. It is fast, scales linearly, and yields robust results across a wide range of inputs. It is, however, strictly heuristic; there is no guarantee of optimality according to formal criteria.

A. *The Algorithm*

The core of the HH technique is the construction of a histogram to represent the data. Given sufficient data points and appropriate bin selection, a histogram can approximate the probability density function of a data set. The algorithm may then examine the histogram rather than the underlying data. This greatly reduces the computational requirements.

The main steps in the HH technique are:

1. Construct a histogram of the data.
2. Identify the largest histogram bin.
3. Identify surrounding histogram bins that exceed a user-specified threshold.
4. Classify the data in these bins as a new cluster.
5. Remove the bins from the histogram.
6. If enough points have been classified into clusters, stop. Otherwise, return to step 2.

Steps 2 and 3 represent cluster identification; steps 4 and 5 represent cluster extraction. Figure 2 illustrates steps 2-5 in one dimension. The full version of the algorithm includes additional error checking to ensure that true clusters are not split apart because the threshold value was poorly selected.

B. *Parameter Selection*

The HH technique requires specification of the number of histogram bins, the bin height threshold during clustering, and the minimum amount of data to classify. The number of

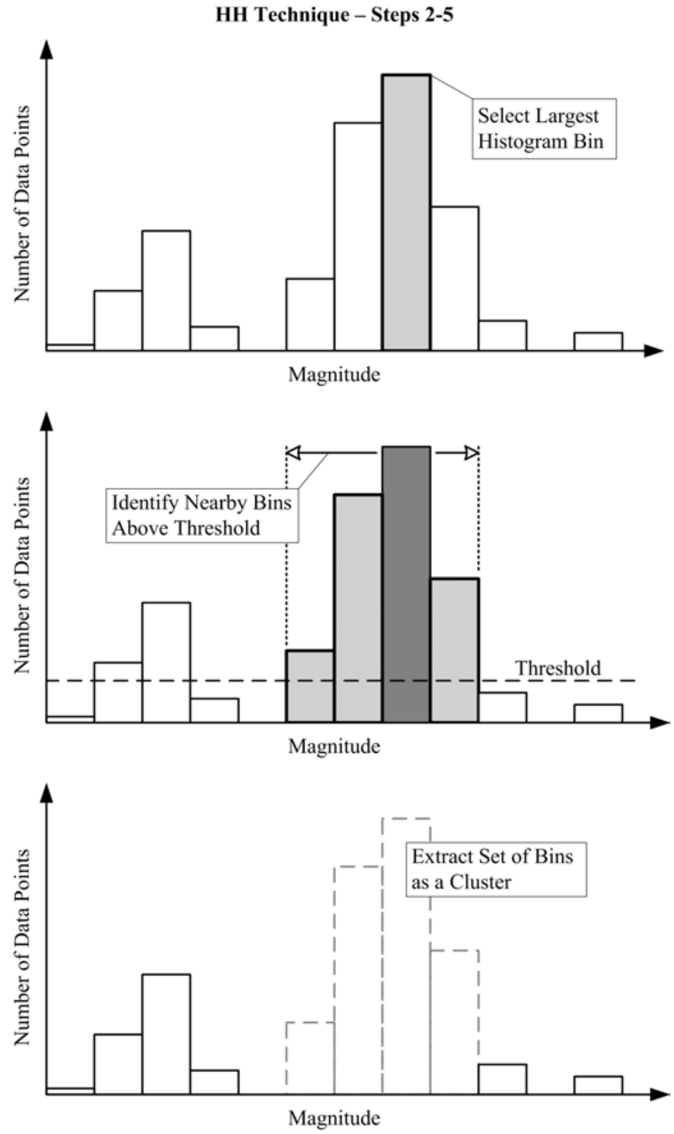

Figure 2. Simplified representation of the histogram heuristic clustering technique in one dimension.

histogram bins drives the resolution at which cluster boundaries may be detected because the width of each bin is inversely proportional to the number of bins used. Generally, more histogram bins are desirable because they enable the algorithm to identify tightly spaced clusters. There are two practical limitations on the number of histogram bins:

- The width of each bin may not be smaller than the precision of the data. Otherwise, some bins will fall entirely between allowable values of the data (and will be empty), artificially dividing the histogram.
- If the number of bins is increased to the point that each bin contains very few points, the histogram loses smoothness and no longer accurately approximates the probability density function.

In practical cases with large data sets, several hundred bins may be used without approaching either limitation.

The bin height threshold controls how the algorithm identifies cluster edges: all adjacent bins above this threshold are included in the cluster. The algorithm uses a relative threshold in order to compensate for the number of bins and adapt to varying cluster densities. If the bin height threshold is too high, the algorithm will truncate the cluster edges. If it is too low, the algorithm will attach bordering noise points to the cluster and may merge multiple clusters together.

The data classification stopping point controls the algorithm by specifying a minimum percentage of the data that must be classified into clusters. This number should be set based on the expected noise level. For example, if the data set contains 1% noise, the data classification threshold should be set at 99% or lower. In general, setting this threshold too low will not affect cluster discovery unless one or more valid clusters contain only a very small number of data points. Setting this threshold too high has an immediate and obvious effect: the algorithm identifies many additional small clusters of noise points.

C. *Advantages and Disadvantages*

The HH technique's primary advantages are simplicity and speed. The algorithm requires no assumptions about the data structure, but still forms intuitive clusters. The bulk of the computation occurs during construction of the histogram, so the computational cost of the HH technique scales linearly. Finally, the HH technique offers reliable noise rejection.

The HH technique has some limitations:

- It requires sufficient data to create a smooth histogram. (This limitation does not usually apply to electrical load data sets, which are typically large.)
- It is sensitive to the initialization of control parameters; however, it is not difficult to select default parameters that function well across a wide range of input data sets.
- It rarely identifies cluster edges perfectly.
- It is unsuited to data sets with adjacent or overlapping clusters because the density-based clustering process cannot distinguish between them.
- It does not extend readily to data sets of high dimensionality. Increasing the dimensionality by one order decreases the number of data points per bin (reducing histogram smoothness), squares the number of bins required to produce sharp cluster edges, and doubles the number of search directions (increasing computational cost). Therefore, in practice the HH technique is limited to clustering in only a few dimensions.

V. Performance Evaluation

To evaluate the HH technique's performance with respect to electrical load data, we used the R environment for statistical computing [22] to implement and test it. During controlled testing, we compared the HH techniques with three other clustering methods: K-means, EM-based normal mixture modeling (MClust) [23], and DBSCAN [24]. We also applied the HH technique to 163 sets of recorded MELs data to gauge performance with real data.

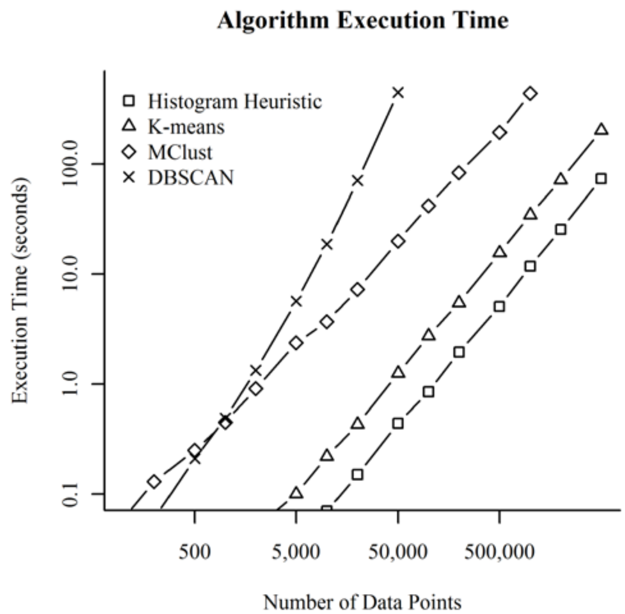

Figure 3. Comparison of histogram heuristic execution time with three other clustering algorithms.

A. *Performance Testing*

During controlled testing, we performed three distinct tests: algorithm sensitivity, quality, and speed. The results of the sensitivity test determined the default algorithm parameters used in the remaining tests.

1) Sensitivity: This test measured the sensitivity of the HH algorithm results to changes in the control parameters. Preliminary testing showed that the data classification stopping point has little effect on the clustering results as long as it is set below the level of the noise. In the test we varied the number of bins and the bin height threshold while keeping the data classification threshold constant at 99%. The test showed that the HH algorithm is robust across a wide range of input parameters. The number of bins may vary from approximately 30 to 3000 and the data classification threshold from approximately 0.1% to 10% without significantly altering the clustering results. Furthermore, the domain of input parameters corresponding to consistent results varied little among differing MELs data sets. Default values exist for these input parameters that perform well with diverse data: we selected 250 histogram bins and a threshold level of 1.0% for the speed and quality tests.

2) Speed: The speed test evaluated execution time of each algorithm using an identical set of control data. The HH algorithm is the fastest of those tested due to its simplicity (see Figure 3). HH algorithm execution time scales linearly with data set size.

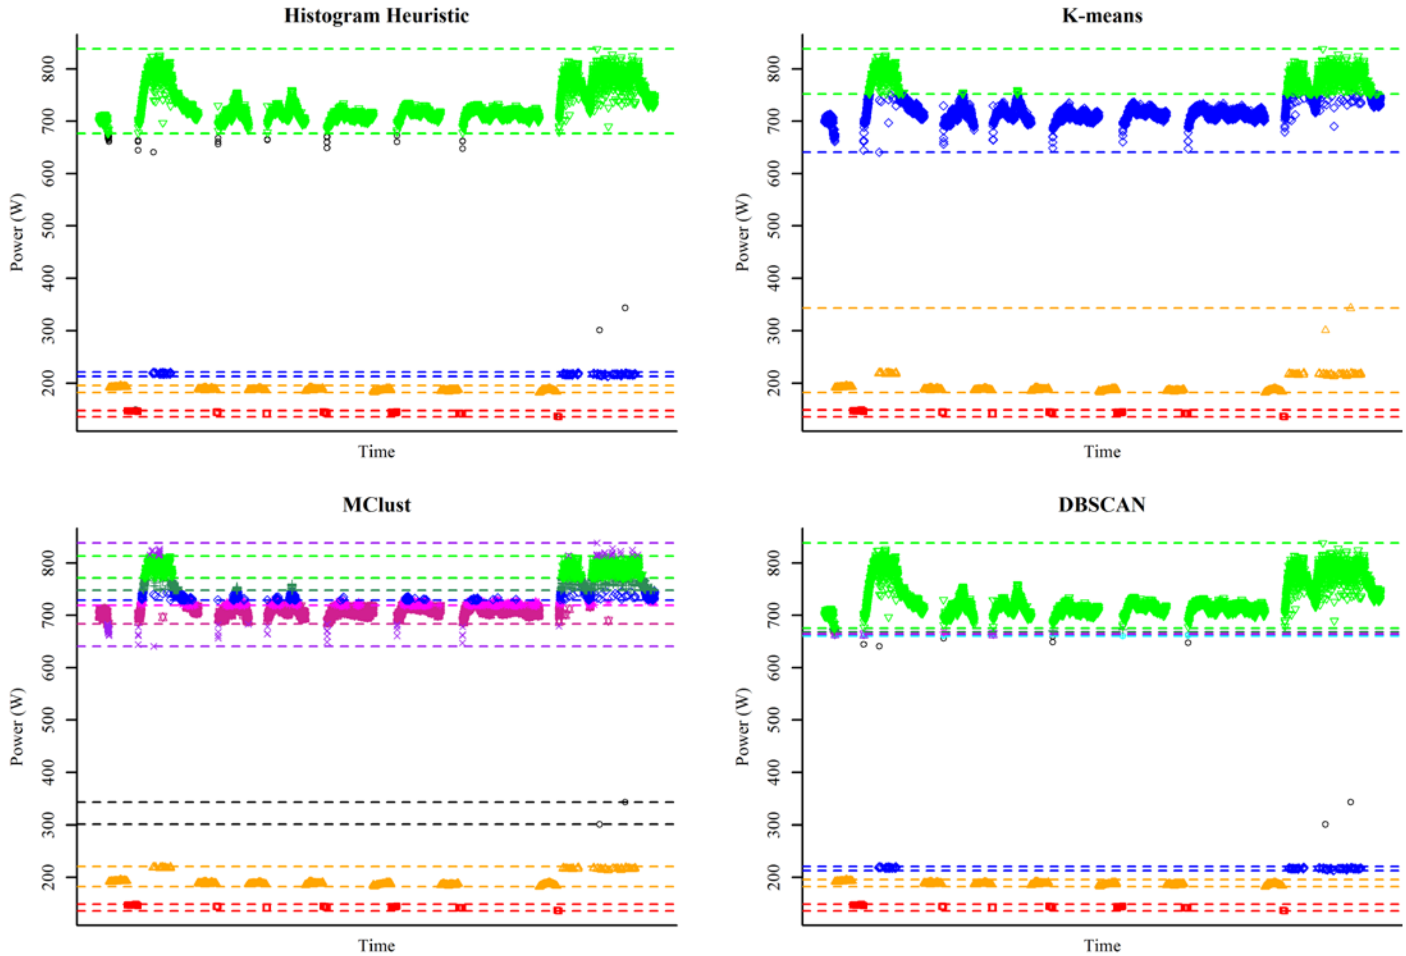

Figure 4. Comparison of clustering results for vending machine load data returned from four clustering algorithms: Histogram Heuristic, K-means, MClust, and DBSCAN.

3) Quality: The quality test compared the quality of output among the four clustering algorithms. Figure 4 displays a qualitative comparison of the four algorithms as applied to the vending machine data of Figure 1. K-means and MClust artificially split the uppermost mode and combine two of the lower modes; only the HH and DBSCAN algorithms return groupings that closely match the inuitive visual power modes. However, HH fails to correctly capture the lower bound of the highest power mode and DBSCAN artificially introduces several small clusters at that bound.

In order to quantify clustering accuracy numerically, we constructed a set of control data with three distinct modes (clusters) plus 1% white noise. (The data was constructed to resemble the observed behavior of a beverage refrigerator.) We initialized each algorithm according to recommended practice, including specifying the number of clusters for the K-means algorithm. We then calculated accuracy by matching indentified and true clusters via best matching of cluster centers and computing the percentage of correctly clustered points. Table I summarizes the results. The HH and DBSCAN algorithms both achieved excellent accuracy. However, DBSCAN identified several small, spurious clusters that the HH algorithm did not (similar to its behavior in Figure 4).

TABLE I. CLUSTERING ACCURACY WITH RESPECT TO CONTROL DATA

Algorithm	Num. Clusters Found	Accuracy
Histogram Heuristic	3 + Noise	99.9%
K-means	3	88.7%
MClust	6[a]	84.9%
DBSCAN	9 + Noise	99.8%

a. MClust classifies noise as an independent cluster

B. *Application: Miscellaneous Electrical Load Data*

We applied the HH algorithm to 163 MELs data sets, plotting each with color-coded modes for rapid visual inspection (not shown in this paper). We inspected the plots, qualitatively assessing the accuracy of the algorithm compared with intuitive groupings. Of the data sets analyzed, 26 were single mode devices and 137 were multimodal. The algorithm properly identified all single-mode devices. In 75% of the multimodal cases, the clusters returned were sufficiently similar to manually identified clusters to be considered accurate. The most common misidentifications (61%) were a failure of the algorithm to capture mode edges accurately (as in Figure 4) or the identification of too many modes. An additional 14% of the errors represent the lumping of two visually distinguishable modes into one because of close or overlapping clusters.

The remaining errors were mostly the result of loads with spiking power use (for example, microwave ovens) for which increased time resolution is needed or mode extraction may not be an appropriate method of characterizing device behavior. In general, the HH technique rapidly and effectively identified the same modes produced by visual inspection. Although a significant number of misidentifications can be corrected by altering the input parameters on a case-by-case basis, doing so is time consuming and demonstrates the need to further refine the default input parameters. In cases where the algorithm performed particularly poorly, visual inspection of the data often revealed either that little or no logical clustering was possible or the presence of some type of data corruption.

In future work, the algorithm could be improved by implementing automatic detection of corrupt data. Other areas for improvement include developing techniques to set the input parameters based on the characteristics of the input data, allowing user-guided searches, and implementing correction by cross-correlation to other data parameters, such as voltage or ambient temperature.

VI. CONCLUSIONS

The HH technique can effectively identify and extract operating modes from electrical load data, performing significantly faster than other, more general clustering algorithms or manual analysis. This allows rapid, automated analysis of electrical load data in preparation for load modeling. The HH technique provides a foundation on which to build a fully automated electrical mode detection system. The technique provides a path for large-scale MELs characterization, a critical step in decreasing the amount of energy used by MELs in buildings.

ACKNOWLEDGMENT

The authors thank C. Lobato of NREL for his assistance with data collection throughout the project. The authors also thank Dr. L. Brackney of NREL for his technical assistance and feedback, including extensively reviewing this paper. Finally, the authors thank Dr. P.K. Sen of Colorado School of Mines for his review of this paper.

REFERENCES

[1] K. McKenney, M. Guernsey, R. Ponoum, and J. Rosenfeld, “Commercial miscellaneous electric loads: Energy consumption characterization and savings potential in 2008 by building type.” TIAX LLC, Lexington, MA, Tech. Rep. D0498, May 2010, p. 224.

[2] D&R International, Ltd., *2009 Buildings Energy Data Book*. U.S. Department of Energy, 2009.

[3] Energy Information Administration, “Annual Energy Outlook 2009, With Projections to 2030.” U.S. Department of Energy, Washington, DC, Govt. Rep. DOE/EIA-0383(2009), March 2009.

[4] A. de Almeida, P. Fonseca, B. Schlomann, N. Feilberg, and C. Ferreira, “Residential monitoring to decrease energy use and carbon emissions in Europe,” in *Proceedings of the 4th International Conference on Energy Efficiency in Domestic Appliances and Lighting*, London, United Kingdom, 2006, vol. 2, pp. 529-544.

[5] L. Moorefield, B. Frazer, and P. Bendt, “Office plug load field monitoring report.” Ecos, Durango, CO, White Paper, Dec. 2008.

[6] C. Cremer et al., “Energy Consumption of Information and Communication Technology (ICT) in Germany up to 2010.” Fraunhofer ISI and Centre for Energy Policy and Economics, Karlsruhe/Zurich, White Paper, Jan. 2003.

[7] J. A. Roberson, C. A. Webber, M. C. McWhinney, R. E. Brown, M. J. Pinckard, and J. F. Busch, “After-hours Power Status of Office Equipment and\line Energy Use of Miscellaneous Plug-Load Equipment.” Lawrence Berkeley National Laboratory, Berkeley, CA, Tech. Rep. LBNL-53729, May 2004.

[8] *Household Electrical Appliances - Measurement of Standby Power*, IEC Standard 62301, 2005.

[9] G. W. Hart, “Nonintrusive appliance load monitoring,” *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870-1891, 1992.

[10] M. Baranski and J. Voss, “Genetic algorithm for pattern detection in NIALM systems,” in *Proceedings of the 2004 IEEE International Conference on Systems, Man and Cybernetics*, vol. 4, pp. 3462-3468.

[11] A. Marchiori and Q. Han, “Using circuit-level power measurements in household energy management systems,” presented at the First ACM Workshop On Embedded Sensing Systems For Energy-Efficiency In Buildings (BuildSys), Berkeley, CA, 2009.

[12] C. Laughman et al., “Power signature analysis,” *IEEE Power and Energy Magazine*, vol. 1, no. 2, pp. 56-63, 2003.

[13] K. D. Lee, S. B. Leeb, L. K. Norford, P. R. Armstrong, J. Holloway, and S. R. Shaw, “Estimation of Variable-Speed-Drive Power Consumption From Harmonic Content,” *IEEE Transactions on Energy Conversion*, vol. 20, pp. 566-574, Sep. 2005.

[14] L. K. Norford, S. B. Leeb, D. Luo, and S. R. Shaw, “Advanced electrical load monitoring: a wealth of information at low cost.” presented at Diagnostics for Commercial Buildings: from Research to Practice, Pacific Energy Institute, San Francisco, CA, 1999.

[15] K. Roth, “Review of techniques for non-intrusive electrical load monitoring,” Seminar, 13-Aug-2010. [Online]. Available: http://eetd-seminars.lbl.gov/seminar/review-techniques-non-intrusive-electrical-load-monitoring

[16] R. Xu and D. Wunsch, *Clustering*. Hoboken, NJ: John Wiley & Sons, 2009.

[17] P. Tan, M. Steinbach, and V. Kumar, *Introduction to Data Mining*, US ed. Reading, MA: Addison-Wesley, 2005.

[18] M. Ester, H. Kriegel, J. Sander, and X. Xu, “A density-based algorithm for discovering clusters in large spatial databases with noise,” in *Proceedings of 2nd International Conference on Knowledge Discovery and Data Mining*, 1996, pp. 226-231.

[19] C. Fraley and A. E. Raftery, “Model-based clustering, discriminant analysis and density estimation,” *Journal of the American Statistical Association*, vol. 97, pp. 611-631, 2002.

[20] J. Sander, M. Ester, H. Kriegel, and X. Xu, “Density-based clustering in spatial databases: The algorithm GDBSCAN and its applications,” *Data Mining and Knowledge Discovery*, vol. 2, pp. 169-194, 1998.

[21] M. Ankerst, M. M. Breunig, H. Kriegel, and J. Sander, “OPTICS: Ordering points to identify the clustering structure,” *SIGMOD Record*, vol. 28, no. 2, pp. 49-60, 1999.

[22] R Development Core Team, *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing, 2010.

[23] C. Fraley and A. E. Raftery, “MCLUST version 3 for R: Normal mixture modeling and model-based clustering.” Dept. of Statistics, University of Washington, Seattle, WA, Tech. Rep. 504, Sep. 2006.

[24] C. Hennig, fpc: Fixed point clusters, clusterwise regression and discriminant plots. [Online.] Available: http://CRAN.R-project.org/package=fpc