World Academy of Science, Engineering and Technology
International Journal of Electrical and Computer Engineering
Vol:6, No:7, 2012

Performance Monitoring of the Refrigeration System with Minimum Set of Sensors

Radek Fisera, Petr Stluka

***Abstract*—**This paper describes a methodology for remote performance monitoring of retail refrigeration systems. The proposed framework starts with monitoring of the whole refrigeration circuit which allows detecting deviations from expected behavior caused by various faults and degradations. The subsequent diagnostics methods drill down deeper in the equipment hierarchy to more specifically determine root causes. An important feature of the proposed concept is that it does not require any additional sensors, and thus, the performance monitoring solution can be deployed at a low installation cost. Moreover only a minimum of contextual information is required, which also substantially reduces time and cost of the deployment process.

***Keywords*—**Condition monitoring, energy baselining, fault detection and diagnostics, commercial refrigeration.

I. Introduction

THE refrigeration system typically consumes more than 50% of the total supermarket energy [1]. Hence it is highly important to operate the refrigeration system at its optimum performance level. Undetected faults or equipment degradations can cause economic losses and potentially violate existing strict regulations regarding the food quality. The reduction of the equipment downtime, service cost and utility cost are the main drivers for on-going research in the refrigeration fault detection and diagnostics area.

There are several approaches how to handle the faults in the system. The simplest but most expensive is to perform corrective actions only in response to equipment failures – fault based corrective maintenance. Smarter and widely used approach is the so-called preventive maintenance. In this case the maintenance is performed regularly in selected time intervals, which are typically based on the equipment manufacturer recommendation. But still, the particular equipment condition is not taken into account. In contrast to that, the condition based maintenance (CBM) aims to trigger the maintenance action at the time when it is necessary, i.e. when there is a clear evidence of deteriorating performance. If the monitoring is done in a systematic way, many "hard" faults can be detected, which would otherwise cause the system to stop functioning. Moreover in case of degradations (slowly evolving "soft" faults) the optimum maintenance schedule can be determined.

Radek Fisera and Petr Stluka are with the Honeywell Automation & Control Solutions (ACS) Laboratories, Prague, Czech Republic (e-mail: radek.fisera@honeywell.com, petr.stluka@honeywell.com).

Faults can be detected at various levels of the equipment hierarchy: the system or circuit level is good for detecting general faults, while the equipment or device levels can help to detect specific problems, such as the stuck expansion valve. Performance monitoring of the refrigeration system may include aspects of monitoring and processing of alarms, monitoring of process data (temperatures, pressures), and monitoring of electricity consumption.

Typically the refrigeration monitoring and control system activates alarms whenever the measured values of key parameters (e.g. case temperatures, compressor discharge pressure) are out of their predefined ranges. This is the commonly used method for indication and alerting of potential problems. Analysis of alarm logs can provide additional insights. Both manual and advanced pattern recognition methods were described in the literature [2], [3]. Sequences or combinations of alarms can be learned from historical data, and consequently used in real-time for detection of specific faults in the current operation.

From the on-line monitored process data it is possible to calculate performance metrics characterizing the system as a whole, or its individual parts. Coefficient of performance (COP), which is widely used for monitoring of chillers and other vapour compression cycle equipment, might be used as a natural metric for the system-level monitoring. COP is evaluated as the ratio between the delivered cooling energy and the total energy input. However the coefficient is very hard to calculate due to usually missing mass flow measurements. Cooling output is then often replaced by the cooling demand, assuming that it is met, calculated by a model, which takes into account occupancy schedules, case door open signals and other inputs.

In the electricity consumption monitoring scenario, the measured power is systematically compared with a referential value (baseline) and any major discrepancy is reported. Several approaches can be used to construct the baseline. First one determines the expected energy consumption from the manufacturer data for electrical devices (compressors, fans, door heaters) adjusted to the actual operating conditions [4]. The other way is to model the baseline statistically as a dependency between the energy consumption and suitable explanatory variables.

Independently of its technical approach, any performance monitoring solution has to address the following list of typical requirements and desired features, which is a blend of end-user's and solution provider's perspectives..

- Simplicity and robustness – ensure wide solution applicability without additional tuning work.
- No additional sensors – this requirement is a typical barrier for the deployment of more sophisticated methods that require more data.
- Minimum deployment efforts (costs) – the less context information is required the better.
- Reasonable computational burden – applies to both controller-embedded (on-line) and server-based (off-line) analytics.
- Model adaptability – as the control strategies are often changed throughout the season (year) to cope with varying operating conditions, any models used in performance monitoring need to be adaptable as well.
- Ability to cope with data inconsistencies, such as missing data or outliers.
- Mitigation of false alarms using confidence bounds or similar techniques.
- Monetization – ability to convert information about faults into cost impact.
- Clearly arranged, easily readable and interpretable result visualization.

This paper describes overall concept of a remote performance monitoring solution whose core functionality is based on a substantially enhanced method for energy monitoring and baselining. An important aspect of the proposed solution is that it does not require more than the "typically available sensor set", which is understood as a group of sensors that can be found practically on any site, despite the rather large variety in types of systems and solution providers involved.

The paper is structured as follows. Firstly the energy baselining methodology is described (section II) including overall system architecture, associated fault and degradation detection capabilities in dedicated subsections. This is followed by examples of lower level monitoring and diagnostics methods in section III. Finally the conclusions are made.

II. Energy Baselining

Several alternative approaches can be used for development of energy baseline models for commercial refrigeration systems. But each has some pros and cons.

Energy models can be based on the first principles: physics, thermodynamics or chemistry. But they require a lot of contextual information that is very difficult to get in an automated way, because there is no standard common language (ontology) used by the refrigeration vendors or solution providers for the system description. So any model which requires rich contextual information is not only quite difficult and time consuming to build but it is also hardly portable to other systems.

Statistical models bring more robustness in terms of portability and scalability as they are often based only on basic measurements and relationships in the refrigeration system while neglecting some details. As the need for contextual information is significantly smaller, these methods are relatively widely applicable, robust and having low deployment costs. On the other hand they may be less accurate when applied to time series data with short sampling period (in minutes). Usually a reasonably good compromise model structure can be found as will be also documented in following sections.

A. Remote Monitoring Architecture

The remote performance monitoring system itself consists of several parts as illustrated in Fig. 1.

System data and electricity measurements are collected by the local control system and transferred via a safe link (TCP/IP based) to the data warehouse, which is deployed in a remote data center. This data transfer is usually performed in a batch-wise manner, e.g. once per hour. This is appropriate as the remote performance monitoring analytics are not supposed to run in real-time. Instead, they are used to support the interactive work of an energy analyst in the remote data center.

The data is further processed in the remote data center. Raw data integrity checks are applied to the selection of data points needed for specific calculations and evident outliers are removed. Consequently, the data points are synchronized and optionally aggregated in time (e.g. averaged) before they enter the core baseline modeling algorithm. The comparison of the measured actual data with the expected energy consumption (provided with confidence bounds) produces a list of deviations, which are further processed in the reasoning layer. The final output for the user (energy analyst or operator in the remote monitoring centre) is the information whether the monitored system is working properly or if there are any obvious faults or issues that require further investigation.

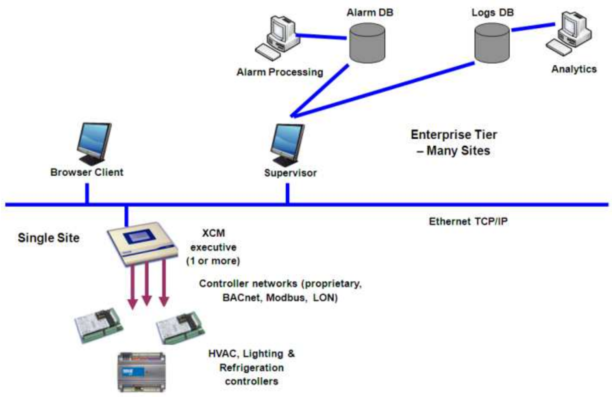

Fig. 1 System Architecture

Two types of faults can be detected by the remote analytics at the refrigeration system (circuit) level – instantaneous anomalies and long-term degradations, which are indicative of faults with completely different dynamics.

B. Data Pre-processing

Invalid or missing data has to be properly handled because otherwise they could cause misleading results. Ideally, some basic form of data pre-processing can be done locally by control hardware, before the data is uploaded to the warehouse. However, this is not always possible.

Three data pre-processing tasks can be distinguished in the remote monitoring system. The first one is to remove raw outliers, i.e. the data outside associated validation ranges. Some less evident outliers can be actually left in the data since the proposed robust regression method (subsection C) is able to cope with them. In the second task, individual time series are synchronized and, in the third one, they are aggregated in time. The aggregation step is important as it helps to mitigate the impact of transient system dynamics. Effective aggregation can eliminate the need for a dynamic baseline model, which would otherwise increase the overall solution complexity.

C. The Core Algorithm

The core algorithm is based on the locally weighted regression technique [5], [6], [7] which was selected as a method allowing to meet many of the challenging requirements summarized in section I.

The model defined in [8] can be adopted for the purpose of energy baseline modeling. Let's assume sequences of independent variables $x_N = (x_1,\ldots,x_N)$ and dependent variable $y_N = (y_1,\ldots,y_N)$, where x_k is a vector of length *m*. Further, it is supposed that the relation between x_k and y_k can be described by stochastic functional relationship $y_k = f(x_k)$, $k = 1,\ldots,N$. The data vector x can be mapped onto a feature vector $\varphi_k=\varphi(x_k)$, possibly of much higher dimension p. This gives a possibility to express the functional mapping $f(.)$ as a parametric model, which in our case is the linear regression for one scalar response y_k.

$$y_k = \varphi_k \theta + \varepsilon_k \tag{1}$$

where ε_k is a noise term and θ is a vector of local model coefficients to be estimated from

$$\theta = \left(\varphi_k{}^T \mathbf{W}_k \varphi_k\right)^{-1} \varphi_k{}^T \mathbf{W}_k \mathbf{y} \tag{2}$$

An estimate of the dependent variable is then obtained by

$$\hat{y}_k = \varphi_k \theta \tag{3}$$

$$\hat{y}_k = \varphi_k \left(\varphi_k{}^T \mathbf{W}_k \varphi_k\right)^{-1} \varphi_k{}^T \mathbf{W}_k \mathbf{y} \tag{4}$$

$$\hat{y}_k = \sum_{i=1}^{N} l_i y_i \tag{5}$$

where $\mathbf{W}_k$ is a weighting matrix. The last equation (5) demonstrates the linearity in *y*.

The refrigeration system is modeled using a selected set of explanatory variables. This is very often some combination of outdoor and indoor air parameters, such as temperature or humidity, augmented with time variables and metrics representing the load imposed by occupants.

The local regression method satisfies the adaptability requirement by including a serial time variable in the set of explanatory variables. Moreover, possible global - in terms of the whole space of operating conditions - non-linear relationships can be effectively approximated by local linear (in coefficients) dependencies.

From the mathematical point of view there is only one strong assumption for calculation of the estimate of the expected energy consumption: the invertibility of the matrix $(\varphi_k{}^T\mathbf{W}_k\varphi_k)$ when computing the weighted least square estimation of regression coefficients according to equation (2). These practical difficulties can be avoided by a thorough data integrity check applied before the estimation algorithm itself.

Application logic, which is adapted to the specific domain of refrigeration systems, divides the required sensor set into two groups – mandatory measurements and optional measurements. The expected energy consumption cannot be calculated in rare cases when any of the mandatory data points is missing. When available, the optional data points help to improve the estimation accuracy that can be observed on the width of calculated confidence bounds.

The local regression algorithm is implemented in a way that a series of predictive models is built on-the-fly for a series of states (query points). These energy consumption models are created locally considering only the most similar data points (N-dimensional list of selected explanatory variables), which are weighted by a selected kernel function. The weighting function secures the localness of the model as it assigns a weight to each data point based on its relative distance, such as normalized Euclidean distance, from the query point in the N-dimensional space. The size of the neighbourhood around the query point, and thereby the number of points used for the model identification, is affected by the chosen kernel weighting function and by the bandwidth parameter applied to each individual explanatory variable.

Though there are methods for an automated bandwidth selection [9], the preferred way is to exploit the domain expert knowledge to estimate optimum values of these parameters in advance. The second important parameter to be chosen is the polynomial order for each explanatory variable. The usual values are 0, 1 or 2. In fact, setting the polynomial order to zero means to build so called nearest neighbours weighted average model [10], [11]. This simple kernel smoothing method is quite popular, however suffers from the large bias especially on the edge of the area given by the kernel function span. Polynomial orders in the presented models - linear, seldom quadratic - were selected according to expected physical dependencies as a trade-off between the bias and variance of the estimate. .

D. Anomaly Detection

Anomaly is a temporal event with unusual energy consumption, i.e. when the energy consumption exceeds the upper confidence bound of the expected value estimate. The expected energy consumption for given time is evaluated using the local data (serial time is also one of explanatory variables) except the data in the nearest time vicinity. This is done in order to exclude potentially anomalous data from the baseline estimation.

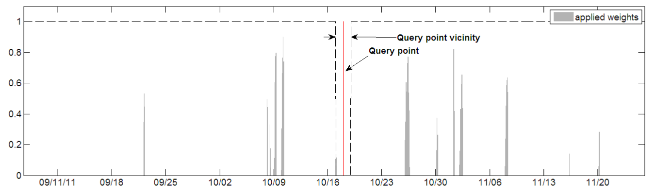

Fig. 2 Weighting for anomaly detection

It actually means that only the data points inside the interval defined by the selected time bandwidth (e.g., 3 months) are used for the baseline evaluation when the kernel with compact support is considered. The same weighting mechanism is applied in all other dimensions. Fig 2 illustrates only the time localness and adaptation capability of the anomaly detection algorithm. Should any set point or even the whole control strategy be changed, the algorithm builds the referential consumption from the similar data within the kernel time span with weights applied to all other explanatory variables – both external and internal conditions. Then the dominant pattern of the similar data points determines the result.

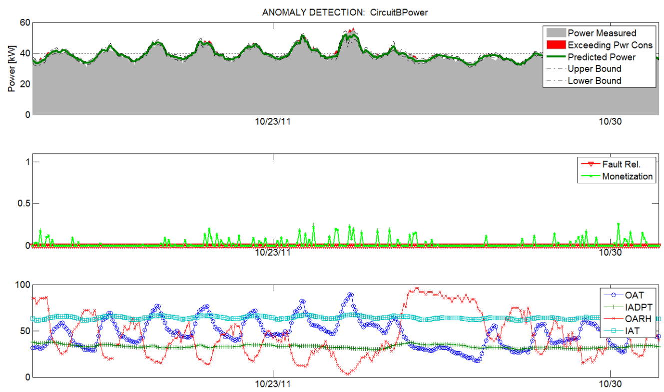

Fig. 3 Anomaly detection example - no anomaly detected

In practice it means that any control strategy change with non-negligible influence is potentially detected only once - possibly as a false alarm - and therefore the energy analyst should check the results against the control strategy data. Anomalies detected immediately after the control strategy has changed should be dismissed.

Fig. 3 provides demonstration of the anomaly detection algorithm applied to real data. The first plot compares the actual energy consumption with the estimated one – i.e. predicted by the local model. The monetization of the fault is straightforward because the discrepancy is measured directly in kilowatts. In this particular example no fault was detected. There were only small deviations that – after processing by the fault reasoning module - didn't flag any specific fault.

E. Degradation Detection

In contrast to anomalies, any degradation is usually a slowly developing process that negatively affects the system performance. Typical examples in the refrigeration system are the refrigerant leakage, condenser coil fouling or compressor oil quality degradation. One implementation challenge of any degradation detection algorithm is that the impact of degradation differs for various system operating conditions. Then the degradation process can be assessed by observing system performance for the same or very similar conditions for a sufficiently long time interval given by the typical degradation dynamics. Mean degradation level over all possible operating conditions for given time span can be then evaluated.

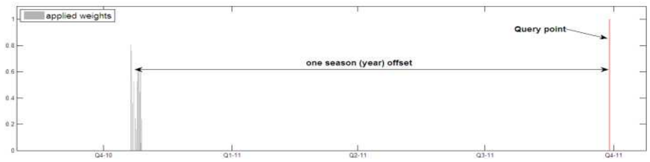

Fig. 4 Weighting function in time dimension - degradation detection

The local models based only on the recent data cannot in principle provide the degradation detection ability because the time distance between the similar data used to build the model and the actually measured data is too short compared to a typical time scale when the degradation has a clear impact on performance. The data used for the baseline construction should be thus drawn from the history when the system operated at its peak performance. The fixed distance between the actually investigated (queried) point and the middle of time weighting kernel function was chosen (Fig. 4).

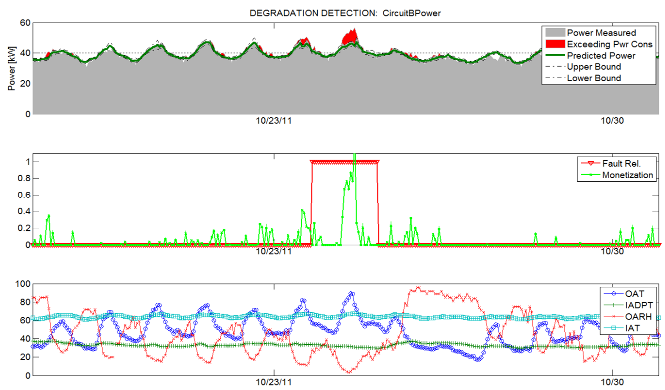

Fig. 5 Degradation detection example - worse performance than last season for certain driving condition was detected

One year proved to be a good choice as the similar conditions in terms of weather can be expected namely in one year old data. The progress of degradation can be assessed on a continuous basis. The increasing trend of the deviation between the measured and expected energy consumption means that the particular equipment or system is deteriorating. Non-negligible variance of this deviation can be observed as the degradation varies for different operating condition throughout the year.

An average degradation trend can be optionally determined from the regressed curve and used as an important input for the optimum maintenance planning. Fig. 5 illustrates the degradation detection algorithm applied to the same data as in Fig. 3. The key difference is that the local models were now built based on the last season data. The results indicate the system was performing better last year for some range of operating conditions. The degradations were evaluated on real test data from several sites using more than one year long history. At some sites the deviation between the current and last year energy consumption for similar operating conditions was significant and helped the analysts at the remote monitoring centre to prioritize further investigations.

F. Fault Reasoning

A clear decision about the fault presence cannot be based only on a single deviation between the actually measured energy consumption and the baseline. Therefore it is important to add a fault reasoning layer at the top of the fault detection functionality. In more general context [12], the system observations specified by one or more rules are called symptoms. Each fault then can have several symptoms that support the particular fault presence (also called contributing or admitting symptoms) and, on the other hand, can also have set of symptoms that deny the particular fault presence (also called excluding or cancelling symptoms). Fault likelihoods can be calculated based on the evaluated relevant (supporting and excluding) symptom values within pre-defined time window.

Two types of faults are considered in the presented concept – anomaly and degradation faults. The relative deviation between measured and upper confidence bound of the estimated baseline is considered as an instantaneous symptom relevancy. Both faults have just one symptom supporting the fault presence and actually no symptom that would deny the fault. Various reasoning techniques [12] can be exploited when transforming symptom to the fault likelihood. The presented solution implements a robust moving average filtering technique. The fault is then reported whenever the aggregated level exceeds a predefined threshold. This substantially reduces the number of events reported to the operator while ensuring that none of the most important is missed.

III. Equipment Level Monitoring

The system level performance indicators discussed in previous sections need to be augmented with equipment level analytics with the capability to detect specific faults and initiate respective corrective actions. Equipment level monitoring significantly narrows the scope of the low level fault searching process. Examples of two analytics applied to compressors - as the most expensive and most critical pieces of equipment within the refrigeration system - are provided in the following sections.

A. Compressor Rack Monitoring

Compressor rack is one of the most vulnerable and expensive parts of the supermarket refrigeration system. Usually there are several compressors (typically 3-6) installed in one rack to deliver the expected cooling load.

The power or amperage of each compressor in the rack is usually available. In this situation, the model for the individual compressor amperage draw can be constructed from available measurements (suction and discharge pressures and temperatures, refrigerant properties and rack control signals). The presented work was focused on the most typical compressor type - reciprocal - however the model structure is general enough so that also other types of compressors can be modeled with sufficient accuracy. When any of compressors in the rack is working inefficiently due to some failure (e.g., leaking valve, increased friction) its capacity is decreased and other compressors have to compensate the missing capacity. This means that the original relations captured in the model are broken and this can be observed as a deviation between expected and modeled compressor amperage for given set of variables.

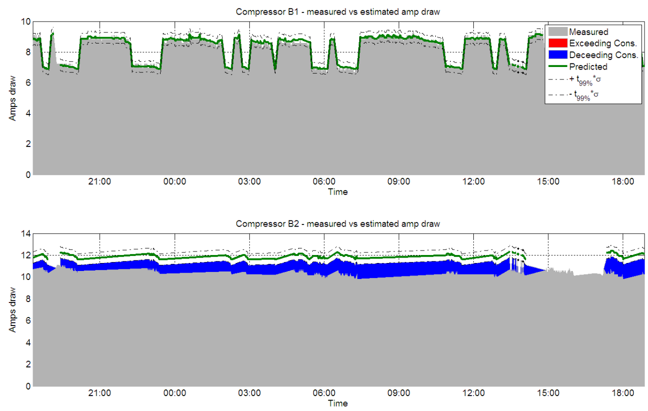

Fig. 6 Compressor amps draw monitoring

Fig. 6 illustrates the effect of an artificially introduced fault on the deviation between the measured amperage and its baseline

B. Liquid Slugging Predictive Detection

Slugging is a process when the large quantities of liquid refrigerant enter the compressor. This can have very detrimental effect on the compressor performance. In some extreme cases, the machine can be completely destroyed. There are several mechanisms, how the liquid slugging originates, described in literature (e.g., [13]). The key task is to predict the liquid slugging before it actually happens, i.e. to detect any dangerous trend in the monitored conditions that will likely lead to the liquid slugging.

One possibility is to exploit existing rule based fault detection frameworks such as described in [12]. In a nutshell, each liquid slugging causing mechanism was described by a set of rules – symptom based on the available measurements. There are number of symptoms that contribute to or cancel the likelihood of the liquid slugging fault. These dependencies are given by a fixed mapping table between faults and symptoms.

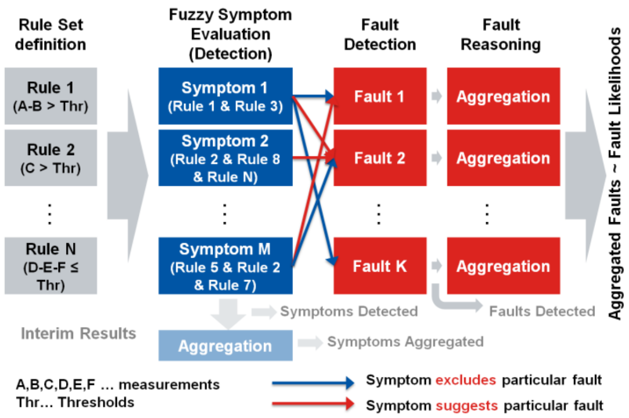

Fig. 7 AFDD engine – example

They are defined on various levels of the refrigeration system according to the particular measurements availability. The methodology is very robust regarding to the sensor requirements. It means that, e.g., the target system can have only the measurements of the refrigerant properties at the common suction line at disposal and the algorithm is still able to provide some useful results.

Of course, a richly instrumented refrigeration system with, e.g. the superheat measurements at each evaporator, can be diagnosed much better, i.e. the liquid slugging fault (or better – increased danger of liquid slugging) can be tracked to its origin. The original approach is enhanced by chaining the AFDD (automated fault detection and diagnostics) engines (see Fig. 7.) from distinct levels of the system.

IV. Conclusion

The paper introduced several concepts and algorithms for remote performance monitoring of the commercial refrigeration systems. In particular, the method for the system (circuit) level relative performance indicator evaluation based on the energy consumption baselining was proposed. It allows distinguishing between the anomalies and degradations detection. Both algorithms were validated against real site test data. Subsequent fault diagnostics can be supported by dedicated equipment level methods described briefly at the end of this paper. Proposed methodology has a very low deployment cost as there are very modest sensor requirements (satisfied already by overwhelming majority of current installations – no additional sensors required) and very limited contextual information is sufficient for the algorithm initialization.

References

[1] EPA, Supermarket Energy Use Profile, Environmental Protection Agency, 2007.
[2] Dan W. Taylor, David W. Corne, David L. Taylor, and Jack Harkness, Predicting Alarms in Supermarket Refrigeration Systems Using Evolved Neural Networks and Evolved Rulesets, In the Proceedings of the World Congress on Computational Intelligence (WCCI-2002), IEEE Press, Honolulu, Hawaii, , 2002, pp. 1988-1993.
[3] Dan W. Taylor and David W. Corne, Refrigerant Leak Prediction in Supermarkets Using Evolved Neural Networks, In the Proceedings of the 4th Asia Pacific Conference on Simulated Evolution and Learning (SEAL), Singapore, 2002.
[4] Abtar Singh, Paul Wickberg, Thomas J Mathews, and Neal Starling, System for remote refrigeration monitoring and diagnostics, US Patent 7644591, 2010.
[5] William S. Cleveland, Robust Locally Weighted Regression and Smoothing Scatterplots, Journal of the American Statistical Association, vol. 74, pp. 829-836, 1979.
[6] William S. Cleveland and Susan J. Devlin, Locally Weighted Regression: An Approach to Regression Analysis by Local Fitting, Journal of the American Statistical Association, vol. 83, pp. 596-610, 1988.
[7] Trevor Hastie, Robert Tibshirani, and Jerome Friedman, The Elements of Statistical Learning, Second Edition: Data Mining, Inference, and Prediction, 2nd ed.: Springer, 2009.
[8] K. Marik, Z. Schindler, and P. Stluka, Decision support tools for advanced energy management, Energy, vol. 33, pp. 858-873, 2008.
[9] Clive R. Loader, Bandwidth selection: classical or plug-in?, The Annals of Statistics, vol. 27, pp. 415-438, 1999.
[10] Geoffrey S. Watson, Smooth regression analysis, Sankhya: The Indian Journal of Statistic, vol. 26, pp. 359-372, 1964.
[11] E. A. Nadaraya, On estimating regression, Theory of Probability and its Applications, vol. 9, pp. 141-142, 1964.
[12] J. Kukal, K. Macek, J. Rojicek, and J. Trojanova, From Symptoms to Faults: Temporal Reasoning Methods, In the Proc. 2009 Int. Conference on Adaptive and Intelligent Systems, Klagenfurt, Austria, 2009.
[13] Danfoss. (2009) Why Compressors Fail - Liquid Slugging. [Online]. http://www.ra.danfoss.com/TechnicalInfo/Approvals/Files/RAPIDFiles/17/Article/LiquidSlugging/Why Compressors Fail Part 3-web.pdf.
[14] Dan W Taylor and David W Corne, An Investigation of the Negative Selection Algorithm for Fault Detection in Refrigeration Systems, In the Proceeding of Second International Conference on Artificial Immune Systems (ICARIS), Edinburgh, UK, 2003, pp. 34-45.
[15] D. Ruppert and M.P.Wand, Multivariate Locally Weighted Least Square Regression, The Annals of Statistics, vol. 22, pp. 1346-1370, 1994.