Data mining and data-driven modelling for Air Handling Unit fault detection

Tianyun Gao^{1,3,*}, Bartosz Boguslawski², Sylvain Marié², Patrick Béguery², Simon Thebault¹, and Stéphane Lecoecue³

¹CSTB, 84 avenue Jean Jaurès, Champs-sur-Marne, 77447 Marne-la-Vallée cedex 2, France

²Schneider Electric, 37 Quai Paul Louis Merlin, 38000 Grenoble, France

³IMT Lille Douai, Department of Informatics and Control Systems, 941 Rue Charles Bourseul, 59500 Douai, France

Abstract. Data-driven automatic fault detection and diagnostics (AFDD) have gained a lot of research attention in recent years. Many existing solutions need to learn from the fault operation data to be able to diagnose the faults. However, these data are usually not available in buildings. In this study we present a data-driven AFDD solution for Air Handling Units (AHUs). The solution consists of three levels of fault detection that require different levels of data availability: the first level is daily energy benchmarking; the second level is control performance evaluation; and the third level is data-driven modelling of mechanical systems. The method is applied to two case studies: experimental data from ASHRAE project 1312-RP, and real-life operation data of an office building in France. These tests show that the solution is able to isolate control faults and mechanical faults of individual components, by learning from normal operation data only.

1 Introduction

Heating, ventilation, and air-conditioning (HVAC) equipment faults and operational errors result in comfort issues and waste of energy in buildings. Thanks to the rapid development of information technologies, sensors, and direct digital controllers, massive amounts of operation data are collected in buildings. However, these data are often not fully exploited to reveal the faults. Efficient data-driven automatic fault detection and diagnostic (AFDD) solutions are of great need in the industry.

In recent years, a number of studies utilized statistical methods, supervised machine learning, and unsupervised machine learning for AFDD in buildings.

Statistical methods analyse the correlation between data and detect the change of correlation when faults occur. For example, Xiao et al [1] applied principle component analysis (PCA) on AHUs to detect faults.

Supervised machine learning methods include regression and classification. Regression methods use normal operation data to train the model, and then compare the prediction with real measurements to detect faults. Wang and Chen [2] applied the auto-regressive exogenous (ARX) model to VAV system. Ajib et al [3] employed a piecewise ARX model to simulate building thermal dynamics. Classification methods use normal operation and abnormal operation data to train the model, and to recognize patterns of different faults. Li et al [4] developed a tree-structured Fault Dependence Kernel (TFDK) method for AHU. Du et al [5] used artificial neural networks (ANNs) to train regression models of AHU, and used the latter to classify detected faults.

Unsupervised machine learning methods are used when no labelled training data are available. Li and Wen [6] developed an AHU fault detection strategy based on pattern matching. Miller, Nagy, and Schlueter [7] developed a day-typing process that uses symbolic aggregate approximation (SAX), motif and discord extraction, and clustering to detect the underlying structure of building performance data.

Although a lot of data-driven AFDD solutions have been developed, most of them have limited capability in diagnostics, unless utilizing fault operation data which are usually not available in buildings.

In this study we present a data-driven AFDD solution consisting of three levels of fault detections. The first and second levels use statistical methods to identify overall abnormal behaviour and control faults, respectively. The third level uses supervised machine learning regression based on normal data to diagnose mechanical faults.

2 Three-level fault detection

An HVAC equipment is usually composed of a) the mechanical system, including heating, cooling and ventilation components, actuators, and sensors; b) the controllers; and c) the human-machine interface (HMI). Components and related faults in an AHU are listed in Table 1 [8].

In practice it is essential to address the faulty component, so that building operators are able to take timely maintenance actions. At the same time, to implement a fault detection solution, the building operators would like to know which data they have to collect, and what results they can expect with these data.

* Corresponding author: tianyun.gao@cstb.fr

Table 1. AHU components and related faults.

Components		Faults
Mechanical system	Cooling and heating coil	Fouling, reduced capacity
	Valves and dampers	Stuck, leakage
	Supply and return fan	Reduced capacity
	Variable speed drive	Not reacting
	Sensors	Drift, offset, stale
Controller		Disabled, overridden, oscillation
HMI		Wrong set-point and wrong schedule

Our three-level AFDD solution is illustrated in Fig 1. The idea of level 1 is to use minimum amount of data to detect abnormal behaviour of the overall system. Then we further provide level-2 and level-3 fault detections, which require more detailed data, to diagnose faults of the controller and the mechanical system, respectively.

Fig. 1. Concept illustration of the three-level AFDD solution.

2.1. Level 1: daily energy benchmarking

Our level-1 AFDD requires the following data:

- Daily energy consumption from sub-meters: cooling coil, heating coil, and fan energy.
- Daily average outside temperature.
- Daily average of controlled variables: supply air temperature and pressure (optional).

In normal operation, the daily energy consumption is correlated with the outdoor air temperature. When the data fall out of the normal pattern, that indicates the occurrence of faults. The ASHRAE guideline 14 [9, Appendix D] describes various linear regression models for building energy consumption. Similar methods are applicable to AHU energy use as well.

Sometimes reduced energy use is caused by reduced comfort in the building. In such cases, the controlled variables, which indicate the comfort achieved by the HVAC equipment, are helpful in diagnosing the abnormal behaviour.

2.2 Level 2: control performance evaluation

Our level-2 AFDD requires the following historical data with system-appropriate sampling time:

- Controlled variables: supply air temperature and pressure.
- Control set points: supply air temperature and pressure set points.
- Control commands: heating and cooling valve commands, supply and return fan commands.

This level uses statistical methods to evaluate control precision and stability.

2.2.1 Precision

Controller minimizes the control error (difference between controlled variable and control set-point) by changing the control command to regulate the mechanical system. A high control error should only occur when the controller is saturated at 0% or 100%. Any other cases indicate overridden, wrong control logic, or slow reaction.

2.2.2 Stability

Without losing precision, overshooting and oscillation of the controlled variable should be minimized. Stability is measured by the cumulated absolute change of the controlled variable.

2.3 Level 3: data-driven modelling of the mechanical system

Our level-3 AFDD requires the following historical data with system-appropriate sampling time:

Heating and cooling coil:

- Outputs: heating/cooling power.
- Inputs: valve command.
- Optional inputs: air entering temperature, air flow rate, and water entering temperature.

Supply and return air fan:

- Outputs: air flow rate.
- Inputs: fan speed command
- Optional inputs: differential pressure over the fan, supply or return air pressure, and fan power.

This level of fault detection trains regression models of the mechanical system with normal operation data to predict the normal output with given inputs. We build four models for heating coil, cooling coil, supply fan, and return fan, and use them to detect faults associated with each component by comparing the measured output with the predicted normal output.

The inputs of the models are selected based on physical knowledge. Such a choice gives better results than using all available measurement or forward-selected variables; see Section 3.2.3 for more details.

The rest of this section will introduce the algorithms for model fitting, cross validation, feature selection, and fault detection.

2.3.1 Regression model

In this study, we use the random forest regression model to make predictions. Random forest [10] is based on decision tree algorithm. CART is a widely used decision tree algorithm for both classification and regression. The algorithm splits the feature space recursively into a set of subspaces, and then fits a simple regression model in each subspace. The splitting stops when no significant gain of regression precision can be obtained.

Random forest is an ensemble algorithm that fits a number of decision trees on randomly selected features and sub-samples of the dataset, and takes the average over these trees for better predictive accuracy and to control over-fitting. A typical number of trees is 100.

2.3.2 Cross validation

Root mean square error (RMSE) is a commonly used metric representing the error of data fitting. Typically, regression algorithms fit the training dataset to minimize RMSE. To avoid overfitting, it is necessary to cross-validate the performance of the model.

We use 5-fold cross validation in this study. The procedure is repeated 5 times, taking different sections as test data. The final RMSE is calculated as the average RMSE of all the folds. The result gives the precision of the regression model. It is also used to set thresholds for fault detection later; see Section 2.3.4.

2.3.3 Feature selection

In the machine learning literature, the predicted variable of the model is called *target*, and the inputs of the model are called *features*. In fitting a regression model, irrelevant features not only increase the computational load, but also make it harder for the algorithm to reveal true correlations between variables. Forward feature selection [11] starts by fitting the model with the single feature that gives the best RMSE, and then iteratively adding features by selecting the best, each time according to RMSE. It stops adding features when no significant improvement of RMSE can be obtained.

This method is purely data-driven. It can improve model prediction performance, but when selected features change, the modelled system boundary changes as well. If the fault to be detected falls out of the system boundary, it will not be detected anymore. Therefore, for the purpose of fault detection, it is suggested to select features mainly based on physical knowledge of the system, and cross-check with data-driven results. More details will be discussed in the case study in Section 3.2.3.

2.3.4 Fault detection

Faults that occur on a specific component cause a major discrepancy between measured and predicted values of that component. The prediction residual is compared to a threshold, which is three times the RMSE obtained from cross validation of the regression model (see Section 2.3.2). When the residual of a specific component exceeds the threshold, it is regarded as a fault.

A fault of some component may cause minor discrepancy on other components, and cause fault alarms on both components. In this case, the component with the larger normalized residual is identified as the real faulty component. Here, “normalized residual” means the residual divided by the fault detection threshold.

3 Case study 1: lab measurement data

3.1. System description

Our first case study used data of the ASHRAE project 1312-RP. Two AHUs were running in parallel in the laboratory: one in normal operation, and one in abnormal

operation with different faults in summer, winter, and spring. Every fault case was run for one day from 6 to 18 o'clock. The AHU was composed of heating coil, cooling coil, supply fan, return fan, and mixing dampers, as shown in Fig 2. The heating and cooling valves were regulated to maintain supply air temperature. The supply fan speed was regulated to maintain supply air pressure. The return fan speed was 80% of the supply fan speed. The economizer switched between full outside-air mode (free cooling) and mixing-air mode (heat recovery) based on the outdoor air temperature. The AHU was serving 12 rooms equipped with Variable Air Volume (VAV) boxes, which regulated the damper individually in each room to maintain room temperature. The collected measurement data are listed in Table 2.

Fig. 2. ASHRAE project 1312-RP, AHU composition.

Table 2. ASHRAE project 1312-RP, description of data.

Short name	Description
Toa	Outdoor air temperature
Tma	Mixed air temperature
Thcda	Heating coil depart air temperature
Thcew	Heating coil enter water temperature
Tccda	Cooling coil depart air temperature
Tccew	Cooling coil enter water temperature
Tsa	Supply air temperature
Tra	Return air temperature
Psa	Supply air pressure
dPsf	Supply fan pressure difference
dPrf	Return fan pressure difference
φsa	Supply air humidity
φra	Return air humidity
Qoa	Outside air flow rate
Qra	Return air flow rate
Qsa	Supply air flow rate
Qhew	Heating coil water flow rate
Qcww	Cooling coil water flow rate
Coadp	Outside air damper command
Cradp	Return air damper command
Ceadp	Exhaust air damper command
Chcv	Heating coil valve command
Cccv	Cooling coil valve command
Csf	Supply fan command
Crff	Return fan command
Ehc	Heating coil heating power
Ecc	Cooling coil cooling power
Wsf	Supply fan electric power
Wrf	Return fan electric power

Fig. 3. Case study 1. ASHRAE project 1312-RP, level 1 fault detection, daily energy consumption

Fig. 4. Case study 1. ASHRAE project 1312-RP, level 2 fault detection, precision analysis

Fig. 5. Case study 1. ASHRAE project 1312-RP, level 2 fault detection, stability analysis

Fig. 6. The impact of AHU mechanical faults on system components

3.2. Results

3.2.1. Level 1: daily energy benchmarking

Daily energy consumption and average supply air temperature and pressure for all normal and abnormal operation cases are plotted in Fig 3. One can see that results of some fault cases significantly drift away from the normal linear pattern. Below is a summary of possible faults.

* Corresponding author: tianyun.gao@cstb.fr

a. Increased heating and cooling energy at the same time. This may indicate heating or cooling valve reversed, stuck open, or having leakage.

b. Reduced cooling energy, increased supply and return fan energy. This may indicate that the cooling valve is stuck closed in cooling seasons. As a result, supply air temperature cannot be maintained at the set-point. All VAV boxes in the rooms require more air for cooling, causing increased fan energy. The supply air pressure may be lower than normal if the supply fan has reached the maximum speed.

c. Increased cooling energy. This may indicate an outside air temperature sensor bias, causing free cooling mode not to be enabled on time.

d. Increased supply and return fan energy. This may indicate coil fouling, dirty filters, stuck-closed air dampers, or other faults which increase the AHU ventilation resistance.

e. Reduced return fan energy. This may indicate stuck closed exhaust damper, return fan failure, or return fan variable speed drive being frozen at low speed.

f. Increased return fan energy. This may indicate return fan variable-speed drive being frozen at high speed.

g. Heating or cooling energy around zero, with no significant increase of supply fan energy. This may indicate stuck closed outside air damper. In this case, if the supply temperature and pressure data are available, they should show that supply temperature is maintained at the set-point, while supply pressure is around zero. The supply fan is running with almost no air flow. It is a critical situation which could damage the fan motor and blow up the duct.

Reduced heating or cooling capacity caused by reduced supply water flow from the pump or pipe leakage does not have influence on daily energy consumption, as long as the maximum capacity is not reached, because it is compensated for by the control. This fault can be detected in our level-3 fault detection. Unstable control does not have influence on daily energy consumption, either. This fault can however be detected in our level-2 fault detection, as we see below.

3.2.2. Level 2: control performance evaluation

Fig 4 shows daily average control errors versus control commands. The green area indicates normal operation. The control errors are close to zero unless the controller is saturated. On this plot we can discover that the “cooling valve stuck” test cases were actually implemented by overriding the cooling command, and not by blocking the valve. Therefore, the dataset is not representative of actual faults. For this reason, before applying the next level of fault detection (level 3), we replace the overridden cooling commands with a ‘normal’ control output based on the real control error.

The daily cumulated absolute change of control error can be used to monitor control stability. As shown in Fig 5, the test case ‘unstable cooling control’ has significantly larger values than the other cases. We hence conclude that there is a stability fault. The test

case ‘unstable fan control’ does not have significant impact, as stated in ASHRAE 1312-RP final report [12]. Test cases are described in table 3.

3.2.3. Level 3: data-driven modelling equipment

Table 3. ASHRAE project 1312-RP, faults associated to each component. ‘*’ indicate not discoverable faults.

Com- ponent	Associated faults	Test cases
Heating coil model	Heating capacity reduced	20,21,22
	Heating coil fouling	18,19
	Heating coil leakage	5,6,7
Cooling coil model	Cooling coil valve stuck closed	4,9,24, 28
	Cooling coil valve stuck open or has leakage	8,10, 23,35,36
	Cooling coil valve reversed	11
Supply fan model	Outside air damper stuck closed	3,25, 29,30
	*Outside air damper stuck midway in mixing air mode	12,13, 26,27
	Heating coil fouling	18,19
	Air duct leak before supply fan	15
	Air duct leak after supply fan	14
Return fan model	Return fan failure	2,33
	Return fan speed frozen	1,37,38
	Exhaust damper stuck closed	17,31,32
	*Exhaust damper stuck open	16,39

Fig 6 illustrates how AHU faults relate to each other and to global system behaviour. For each mechanical component, the correlation between the inputs and the output is dependent on the characteristics of the system. Occurrence of a fault changes the characteristics. As a result, the real output deviates from the prediction.

Based on Fig 6, each fault is associated with one AHU component, as shown in Table 3. The impact of some faults is insignificant [12]. These test cases are marked with ‘*’ in the table. Control faults (test case 14 ‘Cooling coil control unstable’ and test case 42 ‘Fan control unstable’) are to be detected in level-2 fault detection. Outdoor air temperature sensor bias (test cases 40,41) are to be detected in level-1 fault detection.

Fault detection accuracy is calculated based on the confusion matrix, which is shown in Table 4. The accuracy is computed as the sum of true positive and true negative cases divided by the total number of test cases.

Comparison of regression model precision and fault detection accuracy with different feature selections is shown in Fig 7. The method of forward feature selection is described in Section 2.3.3. As can be seen, using all features results in low prediction errors. Forward selection from all features attains even lower prediction errors, because the amount of training data is limited. Using physical knowledge or control commands has higher prediction errors, but attains better overall fault detection accuracy, since there are fewer false negative cases (missed faults).

Fig. 8. ASHRAE project 1312-RP, level 3 fault detection, prediction results of 42 test cases

Fig. 9. ASHRAE project 1312-RP, level 3 fault detection, identify faulty component based on comparison of model residuals.

Table 4. Confusion matrix

	Fault	No fault
Fault detected	True positive	False positive
No fault detected	False negative	True negative

Fig. 7. Case study 1, ASHRAE project 1312-RP, comparison of feature selections for prediction in level 3 fault detection

Detailed prediction results are shown in Fig 8. Heating coil and cooling coil models use only the control command as input. Supply and return fan models use inputs based on physical knowledge as listed in Section 2.3.

In Fig 8, there is only one false negative case, ‘cooling coil stuck at 20%’. The reason is the impact of this fault is too small to be detected. There are several false positive cases. The reasons for the false positive cases are the following:

a. The heating coil model produces false positive when cooling valve is stuck open or has leakage. The reason for this error is that the heating valve opening up to 80% (caused by cooling valve leakage) is not covered in the training data, where heating valve is opened up to at most 35%. The same reason causes cooling coil false positive when heating coil has leakage.

b. When outdoor air damper is stuck closed, the air flow entering the heating and cooling coil is extremely low. This situation is also not covered in the training data. Therefore, the cooling coil model declares false positive in this situation.

c. The return fan model declares several false positives. From Fig 6 we know that the faults of other components have side effects on the return fan.

When more than one component model detects faults, the contributions of each residual were compared, and the component with the largest residual was identified as the faulty component, as shown in Fig 9. As we can see, most test cases are correctly diagnosed;

* Corresponding author: tianyun.gao@cstb.fr

the incorrect ones are cases 3, 19, and 23. The overall fault diagnostics accuracy is 92.8%.

Fig. 10. Case study 2, office building operation data, level 1 fault detection, daily energy consumption.

Fig. 11. Case study 2, office building operation data, level 2 fault detection, precision analysis.

Fig. 12. Case study 2, office building operation data, level 2 fault detection, stability analysis.

4 Case study 2: office building data

4.1. System description

The second case study used real-life operation data of an office building from March to October 2016. Different from the system in the first case study, this building was equipped with a two-pipe system. Depending on the season, the heat pump produced either chilled water or hot water. There was one coil in the AHU for heating and cooling. The unit had an energy recovery unit with an outside air bypass damper. The air system was running in constant flow.

Due to space limitations, only level-1 and level-2 fault detection results are presented in this paper.

4.2. Results

4.2.1 Level 1: daily energy benchmarking

Daily heating, cooling, and fan energy consumptions are shown in Fig. 10. They show similar patterns as in case study 1: The heating and cooling energy is linearly correlated to outdoor air temperature; the fan energy does not change much with outdoor air temperature. The data do not include very low temperature seasons from November to February. Below are our findings:

a. A few cooling energy and fan energy (marked in red circle) values were extremely low. By checking the detailed operation data, we found out that it was caused by missing meter data or early shut down of the AHU.

b. In transition seasons (when outdoor air temperature was between 10 and 20°C), heating energy showed two very different patterns, as marked in green circles. By checking the detailed operation data, we found out that the pump control was changed in June from running continuously to demand control based on valve position. This measure not only saved pump energy, but also reduced heating energy every time the valve was closed. This observation may indicate a small heating valve leakage.

c. The fan energy showed two patterns as marked in green circles: higher in warm seasons, and lower in cold

seasons. It was caused by different pressure set-points, manually changed by the building operator.

4.2.2 Level 2: control performance evaluation

Fig. 11 depicts precision analysis. The green area indicates normal operation. By investigating the points falling out of the green area, we identified some control faults:

a. Sometimes the cooling command was not reacting to temperature control errors. We found out that it only happened when the outside air bypass damper inside the heat recovery unit was controlling supply air temperature at a higher set-point. The aim of the high set-point was to realize a sequence control between free cooling and mechanical cooling. However, it caused the problem that free cooling was disabled when mechanical cooling started to take control of the supply air temperature. A correct solution is to set the free cooling set-point below that of mechanical cooling.

b. Sometimes the fan control precision fell out of the normal pattern because the AHU was manually shut down, or the fan was overridden for a short period of time.

Stability analysis is given in Fig.12. Stability of fan pressure control was good and consistent on most days. The days with bad stability (higher cumulated absolute change) were those when the fan was switched on and off manually. Temperature control stability was significantly worse in cooling seasons than in heating seasons. By checking the detailed operation data, we found out that bad stability always happened in free cooling mode. It turned out that the control of outside air bypass damper in the heat recovery was very unstable.

4.3. Maintenance suggestions

Some of the findings, such as set-point change, control command override, and meter missing data, explain the data patterns we observed. Some other findings require maintenance of the system. We suggest the following actions: a. Test whether or not the heating/cooling valve has leakage. b. Change the control logic of supply air temperature to sequence control between free cooling and mechanical cooling. This can be realized by giving free cooling a lower set-point. c. Change the control parameter of outside air bypass damper to avoid oscillation.

5 Conclusion

In this paper we have presented a three-level AHU AFDD solution. Required data are listed for each fault detection level. The new solution helps the building operator to set up the data collection plan depending on the fault detection focus.

The method was applied to two case studies. The first case study was a laboratory test including normal operation and manually created abnormal operation with many common faults in AHU. For training the models, only normal operation data were used. Abnormal

operation data were used to test the performance of our solution. We have shown that the solution was able to differentiate between control faults and mechanical faults. Mechanical faults of different components were isolated by the level-3 fault detection with an accuracy of 92.8%.

To test the practical usability of the solution, the second case study used real operation data of an office building. We showed that the level-1 and level-2 fault detections can identify some abnormal operations of the system. The level-3 fault detection, which gives more insight on mechanical equipment, is left as future work.

References

1. Xiao, F., Wang, S., Xu, X., Ge, G. (2009). An isolation enhanced PCA method with expert-based multivariate decoupling for sensor FDD in air-conditioning systems. *Applied Thermal Engineering*, 29(4), 712-722.
2. Wang, H., Chen, Y. (2016). A robust fault detection and diagnosis strategy for multiple faults of VAV air handling units. *Energy and Buildings*, 127, 442-451.
3. Ajib, B., Lefteriu, S., Caucheteux, A., Lecoeuche, S. (2017, July). Building thermal modeling using a hybrid system approach. In *20th World Congress The International Federation of Automatic Control*.
4. Li, D., Zhou, Y., Hu, G., Spanos, C. J. (2016). Fault detection and diagnosis for building cooling system with a tree-structured learning method. *Energy and Buildings*, 127, 540-551.
5. Du, Z., Fan, B., Jin, X., Chi, J. (2014). Fault detection and diagnosis for buildings and HVAC systems using combined neural networks and subtractive clustering analysis. *Building and Environment*, 73, 1-11.
6. Li, S., Wen, J. (2014). Application of pattern matching method for detecting faults in air handling unit system. *Automation in Construction*, 43, 49-58.
7. Miller, C., Nagy, Z., Schlueter, A. (2015). Automated daily pattern filtering of measured building performance data. *Automation in Construction*, 49, 1-17.
8. Hyvarinen, J., Karki, S. (1996). Building optimization and fault diagnosis source book. IEA Annex, 25(1996.8).
9. Haberl, J. S., Claridge, D. E., Culp, C. (2005). ASHRAE's Guideline 14-2002 for Measurement of Energy and Demand Savings: How to Determine what was really saved by the retrofit.
10. Friedman, J., Hastie, T., Tibshirani, R. (2001). The elements of statistical learning (Vol. 1, No. 10). New York: Springer series in statistics.
11. Whitney, A. W. (1971). A direct method of nonparametric measurement selection. *IEEE Transactions on Computers*, 100(9), 1100-1103.
12. Li, S., Wen, J. (2014). Application of pattern matching method for detecting faults in air handling unit system. *Automation in Construction*, 43, 49-58.