Adaptable Fault Identification for Smart Buildings

Anika Schumann, Jer Hayes, Pascal Pompey and Olivier Verscheure
IBM Research
Smarter Cities Technology Centre, Dublin, Ireland

Abstract

Malfunctioning HVAC equipment in commercial buildings wastes between 15% and 30% of energy. Many diagnosis approaches tackle this problem, but they either suffer from a lack of detailed fault information or a lack of adaptability to different buildings and equipment. Clearly, especially in the light of an ever increasing amount of sensor data that is available in heavily metered smart buildings, easily adaptable self learning in-depth diagnosis approaches are needed. This paper addresses the challenges of developing such approaches and describes the contribution artificial intelligence techniques like transfer learning, ontologies, knowledge representation or diagnosis can make in overcoming these challenges.

Motivation

There is an increasing need for automated fault identification tools in buildings. Almost 32% of the total energy consumption in industrialized countries is used for electricity, heating, ventilation, and airconditioning (HVAC) in buildings. This value could be significantly reduced if malfunctioning equippment could be identified quickly and numerous diagnosis approaches tackle this problem (Katipamula and Brambley 2005; Youk et al. 2008). However, these methods are generally highly specialized for specifically targetted fault behaviours (Katipamula and Brambley 2005), and thus limited to the diagnosis of well-understood faults. Due to the specialization of these technologies they are difficult to adapt to different equippment types and thus very costly to deploy.

On the other hand, more easily adaptable statistical techniques are used to detect faults in buildings (Jacoba et al. 2010). However, they only provide very limited fault information. Our vision is to develop a practical adaptable fault identification approach that addresses all of the above problems.

In particular such an approach should have the basic ability to discover characteristics of new faults for which there does not yet exist a diagnosis method. This is especially relevant in the context of smart buildings, i.e., buildings that

have an automatic control system and that are heavily metered. Their rich sensor information could potentially be exploited for a more powerful diagnosis. The problems that need to be tackled are thus (i) how to identify and collect all relevant diagnosis information that might allow the characterization of new fault behavior, (ii) how to reuse this diagnostic information across different equipment and buildings and allow for an easy deployment of the approach, and, finally, (iii) how to efficiently exploit the gathered diagnostic information for a timely and optimal identification of faults? This paper discusses the challenges of developing solutions to each of these problems and lists some Artificial Intelligence methods that are promising starting points for tackling these challenges.

Retrieval of Relevant Diagnostic Information

Smart buildings solutions can typically benefit from vast amounts of information that is available via different sources. A big challenge is to identify the relevant data and to develop the technologies that allow their timely retrieval.

For the task of fault identification in buildings there are at least three main sources of information and related technologies that need to be considered:

1. Metered data: data that comes from sensors in the building. It is often in the form of time series and can be exploited using statistical analysis tools like machine learning or time series analysis and forecasting techniques.

2. User feedback: data that comes from the feedback of occupants. This data is often written in free text and requires the use of natural language processing techniques along with well defined ontologies and taxonomies to be exploitable.

3. Domain expert knowledge: expert knowledge of the building's behavior and functioning, generally represented using rule engines. The acquisition of these rules is often done manually leading to a huge work overhead. Rule engines are usually based on pure logic, fuzzy logic or Bayesian belief networks to make use of the knowledge induced by the rules.

Each of these sources of information requires the application of completely different technologies in order to retrieve and use them. Thus the collection of diagnostic information

could be significantly simplified if it were possible to avoid considering all of the above information sources. However, as the following use case demonstrates, this might limit the diagnostic capability.

Consider the case of a building where it is detected that a boiler is consuming an anomalous amount of gas. This boiler is supposed to provide heat to an office room and is controlled by a temperature sensor that makes sure that the target temperature is reached. This anomalous consumption might have (among others) the following explanations:

1. One of the building's air handling units (AHUs) is not working properly leading to an inefficient transport of heat from the boiler to the targeted room.

2. One of the room's windows has been left open leading to an important loss of heat in the room.

3. The temperature sensor is not working properly leading to a wrong report about the room's actual temperature.

Thus, none of the different sources of information listed above can provide a full picture of the observed system:

1. Metered data is insufficient as it cannot capture all that happens in the building. For example, instrumenting each window of the building to know whether it is closed or to what degree it is open is too costly. Furthermore, some variables like the user's thermal sensation are just impossible to measure directly. As it is not feasible to meter everything in the building, metered data provides a partial assertion of the building's situation only.

2. User feedback is insufficient. If asked, users can give the information that a window is open; but they will not detect that the boiler is using much more gas to compensate for it. Users can also complain about the temperature in the room, but can not link it to a defective AHU.

3. Domain expert knowledge enables us to make the link between a boiler consuming too much gas and a defective AHU or a window open, but it is very difficult to apply that knowledge without having metered data and user feedback reporting about the state of the building.

It is therefore necessary to cross-correlate those different information sources in order to develop smart systems for fault identification.

However, each of the information technologies (statistical analysis, rule engines, taxonomy based natural processing engines) needed to process data from different sources is very costly to deploy and maintain. Another difficulty arises from the infeasibility to predefine all fault behavior in advance and from the diversity of the fault concept. A fault might be a drift in energy consumption, a sudden drop in water consumption, or a combination of those. Each of these faults concerns different systems and should be triggered using different data analysis methods on different subsets of variables. Currently, the predefinition of methods and variables is an entirely manual process (Seem 2007; Li, Bowers, and Schnier 2010; Yang et al. 2011).

Techniques are needed that allow a user to incrementally explain to the system how it could identify faults by defining which subset of variables and learning methods the machine should use to detect it.

Summary of Challenges

To summarize, the identification and retrieval of diagnostic information is a challenging problem due to

- the need to analyze data from three different sources,
- the need of integrating data from three different sources,
- the infeasibility of predefining all information required to make complete fault identification and
- the diversity of the fault concept when learning new diagnostic information.

An approach capable of overcoming all these challenges has not yet been developed.

Promising Artificial Intelligence Methods

Techniques from the field of machine learning are promising to meet these challenges. For instance, active learning techniques capable of enabling a domain expert to guide the system in its learning could be considered. In particular the work of (Gervasio, Yeh, and Myers 2011) that presents a metalearning approach could serve as a promising starting point for learning how to define the learning problem based on relevant variables and analysis methods. While the experiments of the above approach were conducted on a synthetic domain and relied on exhaustively generated data, we have to tackle the metalearning challenge by resorting to available data only.

Automatic Reusability of Diagnostic Information

The core problem with reusing the diagnostic information is that it is not straightforward to transfer the different sources of information without considerable effort. In most smart building scenarios metered information is the most readily available information source. It is retrievable via a computer-based control system, the Building Management Systems (BMS). An individual BMS typically monitors and controls a large range of different types of equipment. Thus, each BMS has a large set of variables that represent measurements taken by physical sensors. The use of mark up languages to aid interoperability has already begun, e.g., use of SensorML (Liscano and Kazemi 2010) for sensor systems or PMML for statistical models (Guazzelli, Lin, and Jena 2010). These detailed descriptions can allow an operator to understand similarities between the components of different BMS. However, components of the diagnostic information from one building may still be labeled differently in another building. For example, in one system the term 'Back-end temperature' may be used while in another 'boiler return temperature' may be used to describe the same type of measurement.

Much care must also be taken to transfer diagnostic information retrieved from domain experts. This is often in the form of rules and subtle changes to these rules can have large effects. An example of a simple rule is: $(T_{oa} - T_{ma}) > \epsilon_t$. This same rule appears in (Schein et al. 2006) and (Han and Chang 2009). It essentially looks at the difference between outside air and mixed air temperature under a particular operation condition. However, a simple change in the

thresholds for this rule (ϵ_t) could have a large impact on the efficiency of the building. More generally, transferring statistical information is highly challenging. This is due to the source of information that statistical models are based on: metered data are often very specific to their metered system and it is challenging to segment what statistical knowledge might be common across two buildings and what knowledge from the metered data is specific to the observed system and should not be transfered. Even if the information that can be transfered has been identified, defining how to perform the actual transfer is far from trivial: complex multivariate normalization factors might be needed to apply statistical knowledge from a building with 200 people and 1000 m^2 to a similar building with only 10 people and 200 m^2.

Summary of Challenges

The following problems arise in relation of the transfer of the diagnostic information:

- An ontology based on the reports of users will require techniques to extract the terms and relationships while one based on information supplied by expert users must be built using interactive tools.

- The diagnostic information will not be exhaustive and requires the potential to discover new rules and integrate knowledge from other information sources.

- Even the diagnostic information for the same building can become redundant over time as equipment in the building is changed, zones of use change, etc.

- Different buildings will have different physics, i.e. different termo-dynamic properties, different equipment; and so the diagnostic information will require much adjustment to transfer properly. This adjustment will be expensive in terms of time and man-hours (thus requiring an automatic or semi-automatic approach).

Promising Artificial Intelligence Methods

Ideally a method to automatically transfer the diagnostic information is needed and this may require descriptions which are structured and rich enough to allow for a more seamless transfer of the diagnostic information. The major difficulty here is that there are no freely available domain ontologies and there is no one single solution available to create the application ontology. However, there are methodologies that can aid ontology creation: (1) lexico-syntactic patterns to detect hyponymy relations (Hearst 1992); (2) exploiting the internal structure of phrases to derive taxonomic relations; (3) exploiting hierarchical clustering algorithms to automatically derive term hierarchies from text (Grefenstette 1994; Cimiano, Hotho, and Staab 2005). Deploying the axioms and rules in a real-word setting requires automated variable mapping from the ontology onto the BMS. Ontology mapping is the process of finding semantic correspondences between similar elements of different ontologies. The correspondence can be based on lexical similarity, e.g., (Seco, Veale, and Hayes 2004), among other methods.

A domain ontology for smart buildings must also account for the types of knowledge from all the information

Table 1: Ontological layers for a smart building scenario.

$\forall$x (chilling(x) V heating(x) = fault(x))	AXIOMS and RULES
cool(dom:CHILLER-COIL, range: WATER)	RELATIONS
is-a (FUEL-LINE, PIPE)	TAXONOMY
FAULT := < I,E,L >	CONCEPTS
(regulator, control, controller)	SYNONYMS
HVAC, boiler, chiller	TERMS

sources. Thus in the building management domain there is a great deal of formal knowledge already represented, e.g., in fault detection and diagnostic rules. The ontology builder must reconcile the existing formal knowledge with knowledge from the users (expert or otherwise). Table 1 shows an example of ontological layers for the building management area based on a non-building management (Buitelaar and Magnini 2005). Rather than building an ontology just from the base (term level) an integration with a top-down approach is needed that can exploit existing expert knowledge like rules and axioms.

For tackling the problem of transferring statistical knowledge methods from the field of transfer learning, in particular the work of (Hu, Zheng, and Yang 2010), provide a promising starting point.

Optimal Diagnosis based on iteratively acquired diagnosis information

Given the knowledge that becomes iteratively available there are many different ways of representing it. A major aim is to ensure that the resulting representation allows for easy updates and efficient diagnostic reasoning. Numerous diagnosis approaches for fault identification in buildings have been developed but they are highly specialized for particular fault behaviors (Katipamula and Brambley 2005). Thus, a good understanding of the anticipated fault is needed for selecting an efficient diagnosis approach. In our context this assumption does not necessarily hold. When discovering new fault behaviors one cannot resort to existing specialized diagnosis methods. Therefore, diagnosis algorithms are needed that most efficiently exploit the available knowledge about fault behaviors.

The latter is defined over events from different sources consisting for instance of:

- user feedback, e.g., 'lamp broken',

- results from statistical analysis, e.g., 'fault pattern in time series data of illumination sensor', or

- domain expert input, e.g., 'alert if light switch is on and illumination level below 200 lux'.

The same fault could be identifiable based on any of the above information sources or based on a combination of the latter. Furthermore, some information might require a costly action to obtain it, e.g., request user feedback, and these actions might fail. Thus, in contrast to most diagnosis algorithms that aim to identify faults based on a set or sequence of readily available observations, we face the additional challenge here of obtaining the necessary observations at minimal cost.

Summary of Challenges

For tackling the above mentioned problems a unified diagnosis framework is needed that can

- handle missing sensor information and incomplete system models,

- associate action costs to (observable) events, and

- compactly represent diagnostic information.

Such a framework has not yet been developed.

Promising Artificial Intelligence Methods

Methods from knowledge representation, reasoning and diagnosis have the potential to address these challenges. In fact, parts of the above problems have been tackled before, like the one of missing sensor information and incomplete system models (Bonarini and Sassaroli 1997; Chatain and Jard 2004; Zhao and Ouyang 2008), the one of considering action costs (Torta, Dupre, and Anselma 2008) or the one of compactly representing probabilistic models (Darwiche and Marquis 2002; Dechter and Mateescu 2007). However, these approaches were partially applied to toy systems only and significant advances are necessary in order to integrate them and use them for diagnosing complex building systems.

Conclusions

Properly functioning HVAC equipment is the basis for many challenges in the smart living domain, like reduced energy use and increased occupant comfort. Thus, the problem of developing an easily adaptable self learning fault identification approach for smart buildings is one of great practical significance. This paper has described some of the main challenges that need to be tackled in order to approach this development task and has identified a number of artificial intelligence techniques that have the potential of addressing these challenges.

References

Bonarini, A., and Sassaroli, P. 1997. Opportunistic multimodel diagnosis with imperfect models. *Information Science* 103(1):161–185.

Buitelaar, P., and Magnini, B. 2005. Ontology learning from text: An overview. In *In Paul Buitelaar, P., Cimiano, P., Magnini B. (Eds.), Ontology Learning from Text: Methods, Applications and Evaluation*, 3–12. IOS Press.

Chatain, T., and Jard, C. 2004. Symbolic diagnosis of partially observable concurrent systems. In *FORTE*, 326–342.

Cimiano, P.; Hotho, A.; and Staab, S. 2005. Learning concept hierarchies from text corpora using formal concept analysis. *J. Artif. Int. Res.* 24:305–339.

Darwiche, A., and Marquis, P. 2002. A knowledge compilation map. *J. Artif. Int. Res.* 17:229–264.

Dechter, R., and Mateescu, P. 2007. AND/OR search spaces for graphical models. *Artificial Intelligence Journal* 171:73–106.

Gervasio, M.; Yeh, E.; and Myers, K. 2011. Learning to ask the right questions to help a learner learn. In *IUI'11*.

Grefenstette, G. 1994. *Explorations in Automatic Thesaurus Discovery*. Norwell, MA, USA: Kluwer Academic Publishers.

Guazzelli, A.; Lin, W.-C.; and Jena, T. 2010. *PMML in Action: Unleashing the Power of Open Standards for Data Mining and Predictive Analytics.* CreateSpace.

Han, D.-W., and Chang, Y.-S. 2009. Automated fault diagnosis method for a variable air volume air handling unit. In *ICCAS-SICE, 2009*, 2012 –2017.

Hearst, M. A. 1992. Automatic acquisition of hyponyms from large text corpora. In *COLING*, 539–545.

Hu, D. H.; Zheng, V. W.; and Yang, Q. 2010. Cross-domain activity recognition via transfer learning. *Pervasive and Mobile Computing* In Press, Corrected Proof:–.

Jacoba, D.; Dietza, S.; Komharda, S.; Neumanna, C.; and Herkela, S. 2010. Black-box models for fault detection and performance monitoring of buildings. *Journal of Building Performance Simulation* 3(1):53–62.

Katipamula, S., and Brambley, M. 2005. Methods for fault detection, diagnostics, and prognostics for building systems a review, part i. *HVAC&R RESEARCH* 11(1):3–25.

Li, X.; Bowers, C.; and Schnier, T. 2010. Classification of energy consumption in buildings with outlier detection. *IEEE Transactions on Industrial Electronics* 57(11):3639–3644.

Liscano, R., and Kazemi, K. 2010. Integration of component-based frameworks with sensor modelling languages for the sensor web. In *GLOBECOM Workshops (GC Wkshps), 2010 IEEE*, 235 –240.

Schein, J.; Bushby, S. T.; Castro, N. S.; and House, J. M. 2006. A rule-based fault detection method for air handling units. *Energy and Buildings* 38(12):1485 – 1492.

Seco, N.; Veale, T.; and Hayes, J. 2004. An Intrinsic Information Content Metric for Semantic Similarity in WordNet. In *ECAI'04*.

Seem, J. E. 2007. Using intelligent data analysis to detect abnormal energy consumption in buildings. *Energy and Buildings* 39(1):52 – 58.

Torta, G.; Dupre, D.; and Anselma, L. 2008. Hypothesis discrimination with abstractions based on observation and action costs. In *DX'08*, 253–262.

Yang, X.-B.; Jin, X.-Q.; Du, Z.-M.; and Zhu, Y.-H. 2011. A novel model-based fault detection method for temperature sensor using fractal correlation dimension. *Building and Environment* 46(4):970 – 979.

Youk, S.-J.; Yoo, S.-S.; Lee, B.-K.; Kho, J.-H.; and Lee, G. 2008. Development of fault detection system in air handling unit. In *Convergence and Hybrid Information Technology, 2008. ICHIT '08. International Conference on*, 287 –292.

Zhao, X., and Ouyang, D. 2008. Model-based diagnosis of discrete event systems with an incomplete system model. In *ECAI'08*, 189–193.