Semantics-based expert system for fault detection in air handling units

Sebastian Blechmann[1*], Hannah Görigk[1], Rita Streblow[1], Dirk Müller[1]

1 RWTH Aachen University, E.ON Energy Research Center, Institute for Energy Efficient Buildings and Indoor Climate
Mathieustr. 10
52074 Aachen, Germany
**Corresponding author: sebastian.blechmann@eonerc.rwth-aachen.de*

ABSTRACT

Heating, ventilation and air conditioning systems, e.g., air handling units (AHUs), play a key role in modern building energy systems (BES) as they account for a high share of the energy usage, e.g., 30 % within the U.S. commercial building sector. Unfortunately, many faults occur during the operation of AHUs leading to higher energy usage and the need for expert personnel monitoring the according plants. The latter is a time-consuming task and is assisted by fault detection systems nowadays. While there are many different approaches to detect faults in air handling units, this paper presents an approach based on semantic knowledge.
For this approach, an ontology for fault detection and diagnosis (FDD) in AHUs is developed. The FDD ontology describes which components, e.g., temperature sensors or valves, can face what kinds of faults. Hence, the FDD ontology contains generic knowledge and is interoperable and reusable. Furthermore, generic semantic fault rules based on the air handling performance assessment rules (APAR) are developed.
For this generic concept, the actual information about the examined plants, their configuration, sensors, actuators, and controls are stored in a knowledge graph based on existing ontologies, like Brick Schema.
For the evaluation of the fault rules, the configuration of the plants is queried and rules get activated if the necessary components needed for evaluation of the rules are present. Afterwards, the according time series data for rule evaluation are fetched for analysis, following the linked data principle - in this work from a cloud platform. Depending on the analysis result, fault messages are gathered in a list and sorted according to their incidence to allow a certain prioritization for maintenance personnel. The fault notation follows a unified schema allowing personnel to understand and locate the faults fast.
The developed approach is validated using historical data of a real AHU on our premises. We inject faults into the historical data, e.g., drift, stuck, bias, fouling, abnormal, and frozen. The results indicate a high rate of fault detection, e.g., average accuracy of 99.3 %, precision of 97.85, sensitivity of 95.6 %, and Matthews Correlation Coefficient of 98.4 %. Therefore, the approach is considered valid for fault detection in AHUs.
In the future, it is planned to implement more complex faults, e.g., valve needle separated from drive or control faults. Furthermore, the semantic enriched plant configuration shall be used to improve the discovery of fault-symptom relationships.
The presented approach has the big advantage that the knowledge about air handling units and their faults needs to be formulated only once, leveraging semantic knowledge. Therefore, this generic solution can be used for diverse AHU configurations potentially reducing FDD development efforts.

KEYWORDS

Fault detection and diagnosis (FDD), FDD Ontology, Air Handling Unit (AHU), Fault Component Relationship (FCR), Rule-based FDD

1 INTRODUCTION

Building energy systems (BES) are still responsible for a high energy share, e.g., heating, ventilation and air conditioning (HVAC) systems still account for 30 % of the energy usage (Goetzler et al. 2017). Yet, during the operation of BES, on average 245 faults occur per month and building (Crowe et al. 2023). One common approach to detect faults and, where possible, their causes during their operation is the use of fault detection and diagnosis (FDD) systems. A study by Kramer et al. (Kramer et al. 2020) shows median savings of 6 % in building energy usage in using FDD systems in the first year after installation of the FDD system. Knowledge-based FDD approaches use expert knowledge, e.g., rules to detect abnormalities. In current FDD approaches, manufacturers usually use their own proprietary FDD software, data structure, and fault naming (Chen et al. 2021). This data heterogeneity leads to longer application development time and maintenance personnel need more time to get familiar with faults in a BES due to the absence of a standard for FDD report format. This culminates in the fact that some "fault messages can only be interpreted by FDD tool developers" (Chen et al. 2021).

According to Gallaher et al. (Gallaher et al. 2004), insufficient data interoperability costs $15.8 billion per year in the U. S. capital facilities industry. The issue of syntactic interoperability can be solved by the use of protocol-translating gateways or middleware structures, e.g., cloud platforms. But the issue of semantic interoperability remains. One way to overcome this issue is the use of semantic metadata models enabling operational cost reduction (Bergmann et al. 2022), e.g., in FDD applications. There are already metadata models formalizing semantic descriptions in the building energy sector, all focusing on their very own area, e.g., Brick Schema (Brick) (Balaji et al. 2018) modeling components in BES, like HVAC systems, e.g., air handling units (AHUs), lighting, floors, and rooms. Chen et al. (Chen et al. 2021) created a taxonomy to describe the fault nature, fault component, fault type, fault location, and fault equipment to provide a unified description for fault naming in HVAC systems. The authors provide some examples of possible fault-component (FCR) or fault-symptom relationships (FSR). However, the taxonomy itself does not provide distinct FCR or FSR. Furthermore, in contrast to ontologies, taxonomies do not provide formal descriptions in common standardized language, like the resource description framework (RDF) (Cyganiak et al. 2014).

There are already a few ontology-based expert approaches for FDD in BES: During the development of the proposed FDD ontology, a similar, independently developed approach was published by Hwang et al. (Hwang, Akinci, and Berges 2023). While the authors follow the same approach in modeling and relating certain faults to certain Brick components enabling FSR, they do not use unified fault naming conventions, fault types, or fault locations. Furthermore, faults like malfunctioning, stuck, leakage, fouling, and block are modeled for components like valves, coils, pumps, dampers, and fans by the authors. Yet, common sensor faults are not modeled. In another approach by Pruvost et al. (Pruvost, Wilde, and Enge-Rosenblatt 2023), the authors create an ontology, re-using parts from Brick. They use a rule-based approach to detect faults and risks in the operation of BES. They model symptomatic faults, like simultaneous heating and cooling, overheating, or too low temperatures. Yet, they do not model sensor or actuator faults as such and do not model FCR or FSR. To close this gap, we develop a rule-based expert system for FDD in BES using semantic information based on formalized knowledge. For this, we design an FDD ontology, reusing concepts from Brick, to formalize the modeling of faults and their possible relations according to the taxonomy by Chen et al. (Chen et al. 2021). Additionally, we apply the air handling unit performance assessment rules (APAR) (Schein et al. 2006) to detect faults in the operation of AHUs. We inject faulty data in the historical data of a real AHU and develop a fault detection application written in Python to detect these faults based on the formal knowledge in the FDD ontology. The latter is saved in a fault report for all detected faults to assist maintenance personnel in locating and prioritizing faults.

The remainder is structured as follows:

In chapter 2, we summarize faults and FDD systems in BES. In chapter 3, we describe and evaluate existing metadata models and ontologies for BES followed by the proposed FDD ontology. In chapter 4, we present the used FDD metrics and our use case: A rule-based FDD application on the base of the APAR using knowledge from the FDD ontology. Chapter 5 gives an overview of the results including their discussion. Last, we draw conclusions from this work followed by an outlook in chapter 6.

2 FAULTS AND FDD SYSTEMS IN THE BUILDING ENERGY DOMAIN

There are many faults detected in the operation of BES and they account for a high share of energy consumption in buildings (Crowe et al. 2022; Kramer et al. 2020). However, different studies develop their own fault naming convention and FCR.
Crowe et al. (Crowe et al. 2023) analyze FDD data from 317 buildings in regard to faults in the operation of the BES. The authors originally find 1563 uniquely named faults, apply the unified taxonomy from Chen et al. (Chen et al. 2021), and narrow down the number of unique faults to 182. The taxonomy allows the representation of equipment, e.g., AHU, the component location, e.g., supply air, the component type, e.g., temperature sensor, and last the fault nature, e.g., frozen. Furthermore, the taxonomy allows the classification of faults in condition-based (CB), behavior-based (BB), and outcome-based (OB) faults. Frank et. al (Frank et al. 2019) define a CB fault as a fault that leads to an undesired physical condition inside the plant though they do not always cause faulty behavior. One could argue that these kinds of faults are the "root cause" of other types of faults. An example would be a stuck valve. BB faults are defined by their impact on the behavior of the plant. Often, they do not cause the fault but rather show fault "symptoms". A typical BB fault is simultaneous heating and cooling. OB faults are defined by the overall impact on the metrics of the plant, e.g., high cooling water consumption.
Kim et al. (Kim et al. 2021) analyze 26 different studies regarding fault occurrences in buildings and define 18 types of equipment with each up to 8 different types of faults. Leong (Leong 2019) names fault natures similar to the already discussed taxonomy. The author presents typical faults and components in which they could occur.
Comparing the presented studies, it becomes clear that, there are approaches for unified fault naming in HVAC systems but there is no standard yet. Furthermore, on the one hand, there is no complete list of faults available. On the other hand, the heterogeneous naming schema for faults makes it difficult to have an unambiguous mapping of possible faults to components (Kim et al. 2021).

According to Matetić et al. (Matetić et al. 2022), there are four different categories for FDD approaches: knowledge-based, data-based, physics-based, and hybrid.
In this work, we focus on the use of knowledge-based FDD systems and their wide applicability based on expert knowledge, e.g., formulated in rules. A ruleset often used in FDD expert systems in BES is the set of AHU Performance Assessment Rules (APAR) first described by House et al. (House, Vaezi-Nejad, and Whitcomb 2001). They provide 28 expert rules focussing on temperature control in AHUs depending on the operational state of the plant. These generic rules can find a variety of faults while only depending on eleven sensors in the plant.
To express the connection between the APAR, the components needed for their evaluation, and CB and BB faults, we develop an FDD ontology which is described in the following chapter.

3 MODELING OF BES COMPONENTS AND FAULTS

3.1 Existing metadata models and naming schemas

Semantic metadata models such as Brick, SAREF (Daniele, den Hartog, and Roes 2015), and Project Haystack ("Project Haystack" 2014) can improve the interpretability and interoperability of an FDD System, enabling the widespread use of fault detection across different BES. Even though there is a multitude of semantic metadata models in the building

sector, no existing description can represent a whole BES. Furthermore, there is no standardization between the models and the same object may exist with different descriptions in different models (Bergmann et al. 2022).
An often-used model is Brick which, amongst other things, models BES components in HVAC systems and thermal zones (Balaji et al. 2018). With the current Brick release version 1.3, Brick neither models faults of BES components nor supports FCR or FSR. However, an addition to Brick has been proposed recently by Hwang et al. (Hwang, Akinci, and Berges 2023): Fault-Symptom Brick (FSBrick), which tries to close the gap by adding faults and FCR to Brick. For fault naming, the authors follow the naming conventions defined by Chen et al. (Chen et al. 2021) in their unified taxonomy for HVAC faults. FSBrick has been developed in parallel to the work at hand and was published recently. It is a very promising approach, however, it does not yet model sensor faults, fault types, or fault locations.
Another ontology is proposed by Pruvost et al. (Pruvost, Wilde, and Enge-Rosenblatt 2023). Reusing parts of Brick and other metadata models, they create a set of ontologies aggregated as SENSE ontology which covers symptomatic faults but does not model sensor or actuator faults.
The aforementioned taxonomy proposed by Chen et al. (Chen et al. 2021) covers naming conventions for equipment, faults, and fault locations in order to make faults easy to locate and fix. They introduce a fault name convention that is distinct, reproducible, and, if built upon a standardized knowledge base, applicable to other BES. However, in contrast to ontologies, taxonomies prescribe the fault structure and relationships to certain components but do not describe them in a standardized format, like RDF, allowing knowledge inference.
As heterogenous and incomplete fault mapping proves difficult to make different BES and FDD interoperable and easy to use for maintenance personnel we develop our own hierarchical ontology describing the concepts mentioned above. We apply the taxonomy by Chen et al. as the main structure and naming schema while adding information about BES components and their associated faults based on findings of Leong et al. and Kim et al. (Leong 2019; Kim et al. 2021).

3.2 Proposed FDD ontology

Due to the shortcomings of current metadata models and their fault coverage combined with the purpose of formally expressing unified fault nature, fault types, fault location, and FCR, we decided to create an FDD ontology in RDF syntax to express these factors in one ontology.
For describing the relationships and entities of components and equipment, structures of Brick are reused. All components' data points are modeled as Brick Points, which makes internal naming conventions irrelevant as they are distinctly identified by their Brick type. Furthermore, we assign possible FCR (fdd:hasPotentialFault), a link to historical data (ref:hasTimeseriesReference), and the location of the component (fdd:hasSection). Lastly, we create a link to expert rules needing measurements of these components or equipment in order to be executed (fdd:requiresComponent). The last definition is important to determine the number of rules in a ruleset that fulfill all criteria to be used for fault detection in a plant. For possible FSR, we currently relate rules to certain components based on prior work from Schein (Schein 2006) (fdd:affectedBy). The exact modeling of FSR is important and yet out of scope of this paper. Figure 1 shows a modeling example of an AHU supply air temperature sensor.
In future work, it needs to be evaluated how our approach can be aligned with FSBrick to represent FSR in more detail. Whenever possible, the naming of entities and faults is identical to the taxonomy by Chen et al. to avoid ambiguity.
Using the advantages of inferred knowledge in knowledge graphs using RDF syntax, not all relationships have to be added manually. A reasoner can be used to infer relationships and connections automatically if certain information about a component is known. For example, the assignment of a section to a sensor or an FCR is automatically assigned based on the

corresponding Brick component definition. Furthermore, it can be validated which APAR rules are applicable by checking whether the required components are available. When not all requirements are met, the rule cannot be executed for fault detection in the according plant. When all requirements are met, a rule gets marked activated for this plant.
With the presented FDD ontology and its formalized knowledge, it is possible to create a detailed fault report to make fault localization and fixing as time-efficient and easy as possible.

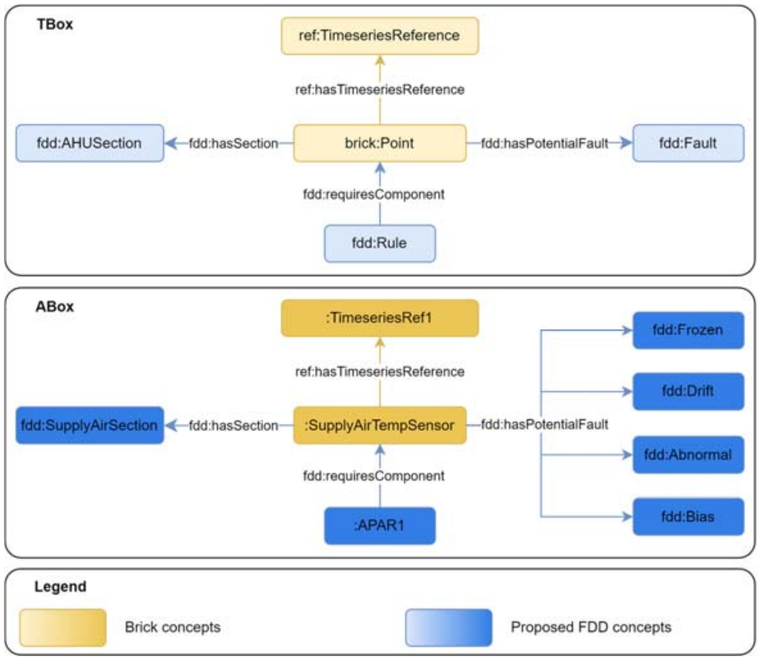

Figure 1: Supply air temperature sensor used in rule 1 of the APAR modeled with the developed FDD ontology.

Disclaimer: The ontology presented corresponds to the status as of 1 July 2024 and is still under development.

4 USE CASE

4.1 FDD application and experiments

The semantic information of a real AHU on our premises, see figure 2, is modeled with Brick and the proposed FDD ontology. The information is stored in a knowledge graph and can be queried from a graph database.
To detect faults in the operation of an AHU, we develop a rule-based FDD application written in Python. The FDD application runs in a loop and checks for possible faults every run. In the FDD application, the APAR set is included. Equation 1 shows rule 1 of the APAR, checking if the supply air temperature (T_{sa}) is smaller than the sum of the mixed air temperature (T_{ma}) and a delta of the supply air fan (T_{sf}) minus a plant-specific tolerance factor (ε_T).

$$T_{sa} < T_{ma} + \Delta T_{sf} - \varepsilon_T \tag{1}$$

To validate the rule, historical data from the according data points of the AHU are evaluated. They are fetched from a cloud platform following the linked data principle.
In each iteration of the FDD application loop with a fixed time period (Δt), in this work 30 minutes, it follows the link towards historical data, fetches it, and evaluates each activated

rule. Before a rule is evaluated, the fetched historical data pass a Hampel filter to find potential outliers.

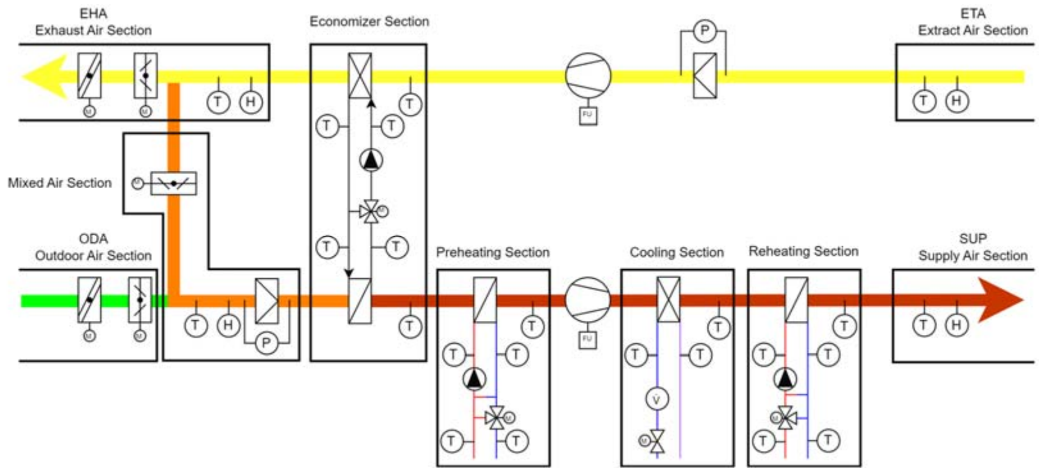

Figure 2: Schematic of used AHU and the proposed sections for fault location assessment.

For validation of the proposed implementation, faulty data is injected into the historical data of the plant before it is fetched by the FDD application. As a reference scenario, untampered historical data is analyzed to exclude possible pre-existing faults from the actual experiments at hand. As proof of concept, a temperature sensor and a valve are chosen as representatives for sensors and actuators, and faults are injected in five experiments at a certain time step (t_{fault}). Experiment 0 is the reference scenario. In experiment 1, drift and stuck faults are injected, in experiment 2, fouling, offset, and frozen faults are injected. In experiment 3, a bias fault is injected while in experiment 4, a combination of drift, stuck, and bias are injected. Last, in experiment 5, a combination of fouling, abnormal, frozen, and bias is injected. The faulty data is injected at the beginning of each experiment (t_{fault}=0) with a data resolution (t) for each experiment of 1 s. All experiments are stopped after a total time frame of 4 h since the AHU operates statically and this work serves just as a proof of concept. This leads to a total amount of data in each experiment of 14,400 per data point.
The actual implementation of each fault is demonstrated in table 1 where each value can be replaced according to the fault description. Depending on the injected fault, the value is replaced by a constant change (e.g., $x_{(t)}$ +/- 2K as offset) or by an increasing or decreasing change dependent on the current time (t) divided by the inspected time period (Δt). The experiments are then evaluated using the metrics from the following section.

Table 1: Fault description for injected faults.

Fault component	**Injected faults**	**Fault description** for $t \geq t_{fault}$	**Fault component**	**Injected fault**	**Fault description** for $t \geq t_{fault}$
	Frozen	$x_{(t)} = x_{(t=tfault)}$		Stuck	$x_{(t)} = x_{(t=tfault)}$
Sensor	Offset	$x_{(t)}$ +/- 2 K	Valve	Fouling	$- x_t * (0.12 * t) / \Delta t$
	Drift	$+/- x_t * (0.12 * t) / \Delta t$			

4.2 Metrics

For the validation of the FDD application, confusion matrices can be used. For that, it can be distinguished between false-positives (FP), false-negatives (FN), true-positives (TP), and true-negatives (TN) predictions over all predictions (n_{pred}). With this, the following metrics are calculated:

$$Accuracy = \frac{TP+TN}{n_{\text{pred}}} \quad (2)$$

$$Precision = \frac{TP}{TP+FP} \quad (3)$$

$$Recall = \frac{TP}{TP+FN} \quad (4)$$

$$F_1 = \frac{2*TP}{2*TP+FP+FN} = \frac{2*Recall*Precision}{Precision+Recall} \quad (5)$$

$$MCC = \frac{TP*TN - FP*FN}{\sqrt{(TP+FP)*(TP+FN)*(TN+FP)*(TN+FN)}} \quad (6)$$

Accuracy, precision, and recall (equations 2-4) are classical metrics in binary classification, living in the range [0, 1]. The commonly used F_1-score (equation 5) ranges in [0, 1] with F1 = 0 standing for all positive samples misclassified and F1 = 1 for perfect classification. We also take the Mathews Correlation Coefficient (MCC) (Matthews 1975), living in the range [-1, 1], into consideration (equation 6). An MCC = 1 represents perfect classification, MCC = 0 represents random classification, and MCC = -1 represents perfect misclassification. The F_1-score does not consider imbalances in the dataset, e.g., if there are way more positive than negative data, in contrast, the MCC does consider such imbalances (Chicco and Jurman 2020) and is therefore taken into account in this study.

For the fault report, we calculate the cumulative sum chart (CUSUM) (Trojanová et al. 2009). The CUSUM represents how often and for how long faults occur. It can be calculated according to equation (7). Here, the fault incidence ($R_i(t)$) rises each time step a fault is detected and decreases every time step the fault is not detected anymore. It is calculated by the maximum of 0 and the binary fault state ($m(f_i,s_j)$) minus a constant (k) plus the previous fault incidence. It is taken as a metric for maintenance personnel to observe and assess detected faults and the duration for which they have persisted; the higher the CUSUM for a fault (i), the longer the according fault persists.

$$R_i(t) = \max(0, m(f_i, s_j) - \mathrm{k} + R_i(t-1)) \quad (7)$$

We use the CUSUM as an indicator for fault incidence which will be displayed in the fault report next to information like the detected fault, fault type, affected component, fault location, faulty and reference values, and the detection time.

5 RESULTS AND DISCUSSION

The metrics indicate a high rate of fault detection, e.g., average accuracy of 99.3 %, precision of 97.85, sensitivity of 95.6 %, and MCC of 98.4 % over all five experiments. Although the FDD algorithm checks for self-injected faults, it does not deliver perfect results. This is due to the fact that some parameters, e.g., the specific tolerance parameter of rule 1, are not chosen perfectly and some faults interfere with each other so that they cannot be isolated and detected easily. The parameters need fine-tuning for each plant. Still, the inference of different faults is not yet properly examined. As mentioned, there are already approaches examining FSR modeling certain BB faults to possible CB faults, e.g., for the APAR (Schein 2006), or alarms to component faults, e.g., based on Brick (FSBrick) that can be modeled and aligned in the FDD ontology in the future. Yet, this is out of scope of this paper and needs thorough discussions of the particular faults. Additionally, further rule sets to detect more complex faults, e.g., a valve needle separated from drive or control faults, can be implemented and analyzed over a longer time frame. Nevertheless, the presented approach is considered a proof of concept for FDD in the presented framework with self-injected faults as the high majority of the injected faults is found based on knowledge derived from the FDD ontology.

As per the fault report, the detected faults are logged in a CSV file and can be displayed in a simple column structure as depicted in table 2. Following the structure of the fault taxonomy by Chen et al., this fault report can distinctly identify faults and aid to locate them quickly. Furthermore, the detected faults can be prioritized based on their incidence, fault type, and faulty values. This is a huge advantage for maintenance personnel since digging through unstructured documents to understand the fault message and locating the actual faulty component are not necessary anymore.

Furthermore, FDD algorithm development engineers can re-use the knowledge from the presented FDD ontology. This saves time and cost in the development of FDD applications. However, the presented metrics only cover the actual FDD and not the portability to other plants. The work at hand delivers a proof of concept of an FDD ontology being used for knowledge-based FDD in a single AHU. Therefore, the approach shall be validated for the use on the other plant configurations in future work, e.g., using a combination of experiments and simulations. This way, fault inferences and the modeling of FSR can actually be examined. In doing so, the idea of the ontology to represent generic, re-usable knowledge can be validated.

Table 2: Exemplary fault report filled with dummy components and values.
(BB = behavior-based, CB = condition-based, i = incidence, n = index)

n	time	component	fault	type	section	i	faulty values
1	10:00	SATsen	Abnormal	BB	SupplyAirSection	0.7	„value“: 18.4, „setpoint“: 21
2	10:00	RHValve	Stuck	CB	ReheatingSection	0.7	„value“: 19, „setpoint“: 28
1	10:30	SATsen	Abnormal	BB	SupplyAirSection	1.4	„value“: 18.5, „setpoint“: 21
2	10:30	RHValve	Stuck	CB	ReheatingSection	1.4	„value“: 19, „setpoint“: 90

6 CONCLUSIONS

Ontologies bear huge potential towards data interoperability and interpretability. Although approaches to standardize semantic metadata schemas in BES exist, faults, their relationships to certain components, as well as fault types and locations are not formally described in ontologies. This is why we presented an FDD ontology and used its information in a rule-based FDD application to detect faults and their characteristics in the operation of an AHU. We injected faults in the historical operational data of the AHU and used the widely applied APAR for the actual FDD. We were able to detect most of the injected faults and create a fault report based on the knowledge modeled in the ontology. This way, maintenance personnel can prioritize faults based on their incidence and type while also locating faults with the given information. Moreover, the ontology allows FDD application development engineers to re-use the concepts presented in this work and adapt them to their use cases. However, the presented ontology is only a small part of the necessary changes to make FDD more usable. For future work, information on how faults affect each other needs thorough discussion and proper modeling, e.g., which BB fault can relate to which CB fault. This is particularly important considering the fact that faults can influence each other making the fault detection of the root cause fault even more difficult. This can be modeled with FSR. Furthermore, the metrics presented in this work only cover the aspect of the actual FDD itself, which is important for

showing the functionality of the approach, but not the portability of the ontology. As the generic knowledge is already formulated, this could be achieved by applying the presented approach to other plants, e.g., with different configurations in future studies reducing time and cost for FDD development. This way the generic sense of this approach can be validated and might open up ways for automatic FDD.
It is planned to publish the proposed FDD ontology with the World Wide Web Consortium.

7 ACKNOWLEDGEMENTS

We gratefully acknowledge the financial support by the German Federal Ministry for Economic Affairs and Climate Action (BMWK), promotional reference 03EN1030B.

8 REFERENCES

Balaji, Bharathan, Arka Bhattacharya, Gabriel Fierro, Jingkun Gao, Joshua Gluck, Dezhi Hong, Aslak Johansen, et al. 2018. “Brick : Metadata Schema for Portable Smart Building Applications.” *Applied Energy* 226 (September):1273–92. https://doi.org/10.1016/j.apenergy.2018.02.091.

Bergmann, Harry, Cory Mosiman, Avijit Saha, Selam Haile, William Livingood, Steve Bushby, Gabe Fierro, et al. 2022. “Semantic Interoperability to Enable Smart, Grid-Interactive Efficient Buildings,” January. https://escholarship.org/uc/item/1325d5j3.

Chen, Yimin, Guanjing Lin, Eliot Crowe, and Jessica Granderson. 2021. “Development of a Unified Taxonomy for HVAC System Faults.” *Energies* 14 (17): 5581. https://doi.org/10.3390/en14175581.

Chicco, Davide, and Giuseppe Jurman. 2020. “The Advantages of the Matthews Correlation Coefficient (MCC) over F1 Score and Accuracy in Binary Classification Evaluation.” *BMC Genomics* 21 (1): 6. https://doi.org/10.1186/s12864-019-6413-7.

Crowe, Eliot, Yimin Chen, Jessica Granderson, Hayden Reeve, Lucas Troup, David Yuill, Yuxuan Chen, and Lawrence Berkeley National Laboratory. 2022. “What We Learned From Analyzing 18 Million Rows of Commercial Buildings’ HVAC Fault Data.” https://escholarship.org/uc/item/8d378066.

Crowe, Eliot, Yimin Chen, Hayden Reeve, David Yuill, Amir Ebrahimifakhar, Yuxuan Chen, Lucas Troup, Amanda Smith, and Jessica Granderson. 2023. “Empirical Analysis of the Prevalence of HVAC Faults in Commercial Buildings.” *Science and Technology for the Built Environment* 29 (10): 1027–38. https://doi.org/10.1080/23744731.2023.2263324.

Cyganiak, Richard, David Wood, Markus Lanthaler, Graham Klyne, Jeremy J. Carroll, and Brian McBride. 2014. “RDF 1.1 Concepts and Abstract Syntax.” February 25, 2014. https://www.w3.org/TR/2014/REC-rdf11-concepts-20140225/.

Daniele, L. M., F. T. H. den Hartog, and J. B. M. Roes. 2015. “Created in Close Interaction with the Industry: The Smart Appliances REFerence (SAREF) Ontology.” *Formal Ontologies Meet Industry : Proceedings 7th International Workshop on Formal Ontologies Meet Industry, FOMI 2015, 5 August 2015, Berlin, Germany*, 100. https://doi.org/10.1007/978-3-319-21545-7_9.

Frank, Stephen, Guanjing Lin, Xin Jin, Rupam Singla, Amanda Farthing, and Jessica Granderson. 2019. “A Performance Evaluation Framework for Building Fault Detection and Diagnosis Algorithms.” *Energy and Buildings* 192 (June):84–92. https://doi.org/10.1016/j.enbuild.2019.03.024.

Gallaher, Michael P., Alan C. O’Connor, John L. Dettbarn, Jr., and Linda T. Gilday. 2004. “Cost Analysis of Inadequate Interoperability in the U.S. Capital Facilities Industry.”

NIST GCR 04-867. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.GCR.04-867.

Goetzler, William, Richard Shandross, Jim Young, Oxana Petritchenko, Decker Ringo, Sam McClive, and Building Technologies Office Corporate. 2017. “Energy Savings Potential and RD&D Opportunities for Commercial Building HVAC Systems.” DOE/EE--1703, 1419622. https://doi.org/10.2172/1419622.

House, John M., Hossein Vaezi-Nejad, and J. Michael Whitcomb. 2001. “An Expert Rule Set for Fault Detection in Air-Handling Units.” *ASHRAE Transactions* 107 (ASHRAE Transactions-2001 Winter Meeting, Atlanta, GA, Volume 107, Part 1): 858–74.

Hwang, Min Young, Burcu Akinci, and Mario Berges. 2023. “FSBrick: An Information Model for Representing Fault-Symptom Relationships in HVAC Systems.” In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, 69–78. BuildSys ’23. Istanbul, Turkey: Association for Computing Machinery. https://doi.org/10.1145/3600100.3623729.

Kim, Janghyun, Kim Trenbath, Jessica Granderson, Yimin Chen, Eliot Crowe, Hayden Reeve, Sarah Newman, and Paul Ehrlich. 2021. “Research Challenges and Directions in HVAC Fault Prevalence.” *Science and Technology for the Built Environment* 27 (5): 624–40. https://doi.org/10.1080/23744731.2021.1898243.

Kramer, Hannah, Guanjing Lin, Claire Curtin, Eliot Crowe, and Jessica Granderson. 2020. “Proving the Business Case for Building Analytics,” October. https://doi.org/10.20357/B7G022.

Leong, Cheng Yew. 2019. “Fault Detection and Diagnosis of Air Handling Unit: A Review.” Edited by Lim Meng Hee. *MATEC Web of Conferences* 255:06001. https://doi.org/10.1051/matecconf/201925506001.

Matetić, Iva, Ivan Štajduhar, Igor Wolf, and Sandi Ljubic. 2022. “A Review of Data-Driven Approaches and Techniques for Fault Detection and Diagnosis in HVAC Systems.” *Sensors* 23 (1): 1. https://doi.org/10.3390/s23010001.

Matthews, B. W. 1975. “Comparison of the Predicted and Observed Secondary Structure of T4 Phage Lysozyme.” *Biochimica et Biophysica Acta (BBA) - Protein Structure* 405 (2): 442–51. https://doi.org/10.1016/0005-2795(75)90109-9.

“Project Haystack.” 2014. Project Haystack Home. May 2014. https://project-haystack.org/.

Pruvost, Hervé, Andreas Wilde, and Olaf Enge-Rosenblatt. 2023. “Ontology-Based Expert System for Automated Monitoring of Building Energy Systems.” *Journal of Computing in Civil Engineering* 37 (1): 04022054. https://doi.org/10.1061/(ASCE)CP.1943-5487.0001065.

Schein, Jeffrey. 2006. “Results from Field Testing of Embedded Air Handling Unit and Variable Air Volume Box Fault Detection Tools.” NIST IR 7365. Gaithersburg, MD: National Institute of Standards and Technology. https://doi.org/10.6028/NIST.IR.7365.

Schein, Jeffrey, Steven T. Bushby, Natascha S. Castro, and John M. House. 2006. “A Rule-Based Fault Detection Method for Air Handling Units.” *Energy and Buildings* 38 (12): 1485–92. https://doi.org/10.1016/j.enbuild.2006.04.014.

Trojanová, J., J. Vass, K. Macek, J. Rojiček, and P. Stluka. 2009. “Fault Diagnosis of Air Handling Units.” *IFAC Proceedings Volumes* 42 (8): 366–71. https://doi.org/10.3182/20090630-4-ES-2003.00061.