



(51) International Patent Classification:
G06Q 50/06 (2012.01)

(21) International Application Number:
PCT/US2017/066687

(22) International Filing Date:
15 December 2017 (15.12.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/435,002 15 December 2016 (15.12.2016) US
62/452,630 31 January 2017 (31.01.2017) US

(71) Applicant: **PRESIDENT AND FELLOWS OF HARVARD COLLEGE** [US/US]; 17 Quincy Street, Cambridge, MA 02138 (US).

(72) Inventors: **YAN, Bin**; 24 Peabody Terrace, Apt. #1302, Cambridge, MA 02138 (US). **SHI, Wenbo**; 13 Ware Street, Apt. 12A, Cambridge, MA 02138 (US). **MALKAWI, Ali M.**; 20 Sumner Road, Cambridge, MA 02138 (US).

(74) Agent: **WEHNER, Daniel T.**; Wolf, Greenfield & Sacks, P.C., 600 Atlantic Avenue, Boston, MA 02210-2206 (US).

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,

DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:
— with international search report (Art. 21(3))

(54) Title: TECHNIQUES OF AUTOMATED FAULT DETECTION AND RELATED SYSTEMS AND METHODS




FIG. 3

(57) Abstract: Methods and apparatus for automatically detecting faults in a building heating ventilation and air conditioning (HVAC) system are described. A fault detection and diagnostic (FDD) framework implements data-driven FDD algorithms to detect faults in system components. The FDD algorithms are configured to use Gaussian process models that consider uncertainties in the operation of the system to reduce false alarms.

TECHNIQUES OF AUTOMATED FAULT DETECTION AND RELATED SYSTEMS AND METHODS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims the benefit under 35 U.S.C. § 119(e) of U.S. Provisional Patent Application No. 62/435,002, filed December 15, 2016, titled “An Open Platform for Automated HVAC Fault Detection and Smart Building Applications,” and U.S. Provisional Patent Application No. 62/452,630, filed January 31, 2017, titled “Automated HVAC Fault Detection Methods and Apparatus,” each of which is hereby incorporated by reference in its entirety.

BACKGROUND

[0001] Commercial buildings consume approximately 19% of all energy consumed in the United States. Heating ventilation and air condition (HVAC) systems often account for approximately 40% of energy consumption in commercial buildings. Faults and deficiencies in control systems are among the most important barriers to energy-efficient buildings. It is reported that around 30% of the energy consumption in buildings is wasted due to system degradation and faults. Potential energy savings on the order of 10 billion U.S. dollars per year may be possible through fault detection and diagnostics (FDD) in large-scale U.S. commercial buildings.

[0002] Forecasting building energy use can be applied to FDD, operation optimization and interactions between buildings and smart grid applications, and has been used in building commissioning to determine retrofit savings. The ability to accurately forecast demand-side loads plays an important role in electric power systems, especially because the future power grid is expected to provide unprecedented flexibility in how energy is generated, distributed and managed. Energy use forecasting is also helpful in fault detection since a baseline prediction is necessary to detect excessive energy consumption.

SUMMARY

[0003] According to some aspects, a computer-implemented method of predicting resource usage within a building is provided, the method comprising receiving data values representing

usage measurements of one or more resources, the data values being produced by one or more sensors located within and/or near the building, identifying a plurality of features, each feature of the plurality of features being a measureable property of one of the one or more resources and thereby associated with the usage measurements of the one of the one or more resources, training, using at least one processor, a first Gaussian process regression model on a first subset of the data values by optimizing first hyperparameters of the first Gaussian process regression model, wherein each of the first hyperparameters is associated with one of the plurality of features, calculating, using the at least one processor, based on the first Gaussian process regression model, an indicator value for each feature of the plurality of features, the indicator value indicating an extent to which variation in values of the feature correlates to variation in the usage measurements produced by the one or more sensors and associated with the feature, training, using the at least one processor, a second Gaussian process regression model on a second subset of the data values by optimizing second hyperparameters of the second Gaussian process regression model, wherein each of the second hyperparameters is associated with one of a subset of the plurality of features, wherein features of the subset are selected from amongst the plurality of features based, at least in part, on said indicator values calculated for the features of the subset, and predicting, using the at least one processor, usage of a first resource of the one or more resources based on the trained second Gaussian process regression model.

[0004] According to some aspects, at least one computer-readable medium is provided comprising instructions that, when executed by at least one processor, perform a method of predicting resource usage within a building, the method comprising receiving data values representing usage measurements of one or more resources, the data values being produced by one or more sensors located within and/or near the building, identifying a plurality of features, each feature of the plurality of features being a measureable property of one of the one or more resources and thereby associated with the usage measurements of the one of the one or more resources, training, using the at least one processor, a first Gaussian process regression model on a first subset of the data values by optimizing first hyperparameters of the first Gaussian process regression model, wherein each of the first hyperparameters is associated with one of the plurality of features, calculating, using the at least one processor, based on the first Gaussian process regression model, an indicator value for each feature of the plurality of features, the indicator value indicating an extent to which variation in values of the feature correlates to variation in the usage measurements produced by the one or more sensors and associated with the feature, training, using the at least one processor, a second Gaussian process regression

model on a second subset of the data values by optimizing second hyperparameters of the second Gaussian process regression model, wherein each of the second hyperparameters is associated with one of a subset of the plurality of features, wherein features of the subset are selected from amongst the plurality of features based, at least in part, on said indicator values calculated for the features of the subset, and predicting, using the at least one processor, usage of a first resource of the one or more resources based on the trained second Gaussian process regression model.

[0005] According to some aspects, a system is provided comprising at least one processor, at least one computer-readable medium comprising instructions that, when executed by the at least one processor, perform a method of predicting resource usage within a building, the method comprising receiving data values representing usage measurements of one or more resources, the data values being produced by one or more sensors located within and/or near the building, identifying a plurality of features, each feature of the plurality of features being a measureable property of one of the one or more resources and thereby associated with the usage measurements of the one of the one or more resources, training, using the at least one processor, a first Gaussian process regression model on a first subset of the data values by optimizing first hyperparameters of the first Gaussian process regression model, wherein each of the first hyperparameters is associated with one of the plurality of features, calculating, using the at least one processor, based on the first Gaussian process regression model, an indicator value for each feature of the plurality of features, the indicator value indicating an extent to which variation in values of the feature correlates to variation in the usage measurements produced by the one or more sensors and associated with the feature, training, using the at least one processor, a second Gaussian process regression model on a second subset of the data values by optimizing second hyperparameters of the second Gaussian process regression model, wherein each of the second hyperparameters is associated with one of a subset of the plurality of features, wherein features of the subset are selected from amongst the plurality of features based, at least in part, on said indicator values calculated for the features of the subset, and predicting, using the at least one processor, usage of a first resource of the one or more resources based on the trained second Gaussian process regression model.

[0006] According to some embodiments, predicting the usage of the first resource comprises predicting a usage amount of the first resource and an estimate of uncertainty associated with the predicted usage amount of the first resource.

[0007] According to some embodiments, the method further comprises detecting an anomaly by comparing additional usage measurements of the first resource with the predicted usage of the first resource.

[0008] According to some embodiments, detecting the anomaly is further based on an estimate of uncertainty associated with the predicted usage of the first resource.

[0009] According to some embodiments, detecting the anomaly is based on identifying that the additional usage measurements of the first resource exceed a threshold confidence interval defined by the estimate of uncertainty associated with the predicted usage of the first resource.

[0010] According to some embodiments, the one or more resources includes one or more of electricity, water, steam, oil and/or gas.

[0011] According to some embodiments, the plurality of features include a first feature that is a measureable property of the first resource, and the first feature indicates usage of the first resource by one of day of the year, week of the year, or day of the week.

[0012] According to some embodiments, the first hyperparameters are characteristic length scales in covariance functions of the first Gaussian process regression model each associated with the one of the plurality of features.

[0013] According to some embodiments, the method further comprises predicting, using the at least one processor, usage of a second resource of the one or more resources based on the trained second Gaussian process regression model.

[0014] According to some embodiments, the method further comprises displaying the predicted usage amount of the first resource to a user via at least one display.

[0015] According to some embodiments, the method further comprises calculating, using the at least one processor, an impact factor of the first resource based at least in part on the predicted usage amount of the first resource and the estimate of uncertainty associated with the predicted usage amount of the first resource.

[0016] According to some embodiments, calculating the impact factor of utilization of the first resource is further based on an amount of noise added to the estimate of uncertainty associated with the predicted usage amount of the first resource.

[0017] According to some aspects of the technology described herein, an open platform for automated HVAC fault detection and smart building applications is described. The platform

may be connected to existing building automation systems using BACnet as well as new smart sensors/devices to monitor and visualize the building and system performance.

[0018] According to some aspects, the platform provides fault detection and diagnostics (FDD) functionality for building managers and facilities to find faults in the HVAC system and provide insights on the potential cost savings.

[0019] According to some aspects, the FDD functionality is implemented by one or more data-driven algorithms that take uncertainty in operations into account.

[0020] According to some aspects, the FDD functionality provides for detecting faults on the component level of an HVAC system and other mechanical system operations.

[0021] According to some aspects, the FDD functionality enables an estimation of energy savings potential using what-if scenarios for building commissioning and retrofitting applications. In some implementations, the what-if scenarios may be executed using actual building data rather than simulation data.

[0022] According to some aspects, the platform integrates interactive and distributed control strategies for optimizing the building performance.

[0023] According to some aspects, the platform is an open platform configured to use web service and Internet-of-things (IoT) technologies.

[0024] According to some aspects, the platform is implemented using a decentralized architecture to address challenges in smart buildings management.

[0025] According to some aspects, the platform integrates one or more advanced FDD algorithms.

[0026] It should be appreciated that all combinations of the foregoing concepts and additional concepts discussed in greater detail below (provided such concepts are not mutually inconsistent) are contemplated as being part of the inventive subject matter disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0027] Various non-limiting embodiments of the technology will be described with reference to the following figures. It should be appreciated that the figures are not necessarily drawn to scale.

[0028] FIG. 1 illustrates a flowchart of some FDD methods for use with building energy

systems;

[0029] FIG. 2 illustrates a schematic of a building commissioning process;

[0030] FIG. 3 illustrates a flowchart of a Bayesian FDD process that may be used in accordance with some embodiments;

[0031] FIG. 4 illustrates a flowchart of a Gaussian process method that may be used in accordance with some embodiments;

[0032] FIG. 5 illustrates a portion of a user interface for enabling a user to select features for use with an FDD algorithm in accordance with some embodiments;

[0033] FIG. 6 illustrates a portion of a user interface for enabling a user to train a model used with an FDD algorithm in accordance with some embodiments;

[0034] FIG. 7 illustrates a portion of a user interface for enabling a user to make predictions using a trained model in accordance with some embodiments;

[0035] FIG. 8 illustrates a portion of a user interface for displaying a timing of building faults detected using an FDD algorithm in accordance with some embodiments;

[0036] FIG. 9 illustrates a flow chart of an illustrative FDD algorithm that may be used to detect faults in a component of an HVAC system in accordance with some embodiments;

[0037] FIG. 10 illustrates a plot of AHU return air temperatures versus outdoor air temperature as used in a parametric uncertainty analysis in accordance with some embodiments;

[0038] FIG. 11 illustrates a bar graph of sub-metered energy consumption of condensing units and AHUs as used in a parametric uncertainty analysis in accordance with some embodiments;

[0039] FIG. 12 illustrates a portion of a user interface for enabling a user to run what-if scenarios using sensor data received by the system.

[0040] FIG. 13 illustrates a schematic of a system architecture within which some embodiments may be used;

[0041] FIG. 14 illustrates a schematic of a system architecture within which some embodiments may be used;

[0042] FIG. 15 illustrates a block diagram of communications in a system architecture within which some embodiments may be used; and

[0043] FIG. 16 illustrates a portion of a user interface for visualizing a whole building energy consumption prediction determined in accordance with some embodiments.

DETAILED DESCRIPTION

[0044] Conventional techniques for detecting faults in a building energy system are labor-intensive and time-consuming. In order to increase the cost-effectiveness, automated fault detection and diagnostics (FDD) is a highly desirable but emerging technology with limited availability in the market. The inventors have recognized and appreciated that conventional techniques for detecting faults in a building energy system may be improved through the use of a novel FDD framework that includes an open platform integrating sensing data from multiple heterogeneous sources. To this end, some embodiments are directed to data-driven FDD processes that take into consideration uncertainties in the modeling process and provide visualization methods that are targeted specifically for FDD.

Table 1: Building energy system fault summary

Category	Device	Fault
Sensor	Temperature Pressure Flowrate	Frozen/bias
Actuator	Damper Coil Valve	Stuck/leak
Equipment	Fan Pump Chiller	Belt slippery/Failure Refrigerant Leaking/Overcharge Condenser Fouling
Control Logic		Unstable, non-optimal
Design		Equipment/duct size

[0045] As shown in Table 1, typical building system faults include sensor faults, control devices (actuator) faults, and control logic faults. The most common faults typically include outside air temperature sensor bias, outside air damper stuck, outside air damper leakage, cooling coil valve stuck, cooling coil valve control fault, heating coil valve leakage, heating coil fouling, AHU duct leakage, mixed air damper unstable, supply fan control unstable, and air filter area blockage.

[0046] The inventors have recognized and appreciated that through rapid automated fault detection and diagnosis (AFDD), the waste in many whole building energy systems (e.g.,

HVAC) can be largely avoided. Most existing techniques for FDD can be roughly divided into three categories: model based, rule-based, and data-driven, as shown in Fig. 1.

[0047] Model-based techniques are typically generated by physical principle laws governing system behavior such as mass and energy balance in terms of static and dynamic models. Model-based fault detection is typically performed by generating residuals values using the data collected from measurements and control signals. During faulty operation, one or more of the residuals is expected to have a value significantly different from normal behavior. Although model-based techniques generally can explain the dynamic behavior of a system well and perform well as an accurate estimator, model-based techniques typically do not work well for real-time computation due to the calculation burden and engineering efforts for detecting and diagnosing sudden faults.

[0048] Rule-based techniques are founded in *a priori* knowledge of the dynamics of the system and are typically derived from expert knowledge in a troubleshooting method similar to how a technician would analyze the data, using a series of if-then type statements. Rule-based methods are calculation efficient and easy to implement. However, the effects of the rules are hard to guarantee, since the rules are often highly-related with the system characteristics and operation modes. Additionally determining the rules is often difficult and typically requires extensive experiments.

[0049] Data-driven techniques develop building operation prediction models based on historical data using statistics. The forecasted operation situation is then compared with the real operation measurements and if the differences between these two operation cases exceed certain thresholds, a fault may be detected and diagnosed. However, training the models properly typically requires a fault-free system and specific faulty operation data. Here fault-free means no single fault exists in any of the whole building energy systems, and specific faulty operation data means identified fault(s) exist in the whole building energy systems, which are conditions that tend to be very difficult to guarantee and obtain in practice. Therefore, FDD results based on the “fault-free” case are hard to guarantee. A potential difficulty with data-driven techniques is the requirement to estimate prior probabilities of faults. Additionally, data-driven techniques typically require fault signatures, described as hypothetical fault behavioral patterns, and it is unclear if the fault signatures used in the prediction models will be applicable across many different applications. To address at least some of these issues, some embodiments are directed

an adaptive probabilistic based Gaussian process and Bayesian network FDD framework for forecasting building energy demands and detecting faults.

[0050] During building commissioning, the operation of subsystems for the building (e.g., HVAC) are verified to ensure that they operate in accordance with particular requirements. In this process, malfunctioning sensors, valves and dampers can be fixed and the system can operate normally. However, these faults are likely to occur again and cause excessive energy consumption. An example of a building commissioning approach is shown in FIG. 2. Building commissioning typically requires the monitoring of thousands of points in a complex building energy management system and involves time-consuming labor-intensive methods such as energy simulation to detect faults in the system. Recommissioning a building typically requires replication of the same time consuming processes involved in commissioning the building at the outset.

[0051] To at least partially address these limitations, some embodiments automatically detect whether faults have occurred again by building a baseline model for whole building energy consumption and comparing the current performance with the baseline model. The baseline model is a Gaussian process model that considers uncertainties in the operation of the systems to more accurately detect real variations in system performance caused by faults. In this approach, observations from an actual process are compared with the outputs from a baseline model and a fault is indicated when the difference between the model outputs and observations is greater than a threshold. Some embodiments detect the increase in energy consumption due to system faults without sending false alarms when the increase in energy consumption actually lies within the range of the uncertainty. Accordingly, considering uncertainty in baseline predictions is important to build a good model.

[0052] Simulation models based on physical principles are typically not good candidates for baseline prediction. Such models are expensive, as they require a deep understanding of the system and model parameters are difficult to estimate. Moreover, physics-based models usually assume that systems are operating under ideal conditions as opposed to reflecting actual system operations, and therefore do not include uncertainties in their predictions. In some embodiments, a data-driven model and Gaussian process regression is used for baseline modeling. Gaussian process regression is preferable to physics-based models because it is able to predict actual system performance based on historical data in an inexpensive way, and because it can quantify prediction uncertainty in the form of a Gaussian distribution.

[0053] FIG. 3 shows a process for determining a baseline model in accordance with some embodiments. As shown, for an existing building, data 310 collected during normal operations, for example, for the first few months following commissioning, may be used during training. Using that period as a training set, a Gaussian process 320 is used to predict baseline energy consumption assuming normal operations. When it is no longer certain whether faults have reoccurred, the Gaussian process 320 may be used to predict baseline consumption. Then, the baseline consumption and observed energy consumption are input into a Bayes classifier 330 to determine in act 340 whether the observed energy consumption is excessive. If the whole building energy consumption is abnormal, the process continues to act 350 to determine component-level FDD using the trained baseline model for comparison to actual data recorded and processed by the system.

Gaussian process regression

[0054] As discussed above, some embodiments are directed to generating a baseline model for predicting faults in a building energy system using Gaussian process regression. The goal of Gaussian process regression is to find the distribution of a nonlinear function $f(\mathbf{x})$ to underlie data points, each of which is composed of input $\mathbf{x}$ and target y . Then the distribution of $f(\mathbf{x}^*)$ can be used to predict the value of y^* .

[0055] N input vectors $\{\mathbf{x}_i\}_{i=1}^N$ can be denoted by $\mathbf{X}$ and the set of corresponding target values $\{y_i\}_{i=1}^N$ by the vector $\mathbf{y}$. Using the Bayes' theorem, the posterior probability distribution of $f(\mathbf{x})$ is:

$$P(f(\mathbf{x})|\mathbf{y}, \mathbf{X}) = \frac{P(\mathbf{y}|f(\mathbf{x}), \mathbf{X})P(f(\mathbf{x}))}{P(\mathbf{y}|\mathbf{X})} \quad (1)$$

[0056] In a regression problem, $\mathbf{y}|f(\mathbf{x}), \mathbf{X}$, the distribution of the target values given the function $f(\mathbf{x})$ is usually assumed to be Gaussian. The prior $P(f(\mathbf{x}))$ may be placed on the space of functions, without parameterizing $f(\mathbf{x})$. A Gaussian process is specified by a mean function (usually a zero function) and a covariance function $k(\mathbf{x}_i, \mathbf{x}_j)$. In some embodiments, the covariance function is a Gaussian kernel:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \exp \left[-\frac{1}{2} (\mathbf{x}_i - \mathbf{x}_j)^T \mathbf{W}^{-1} (\mathbf{x}_i - \mathbf{x}_j) \right] \quad (2),$$

where

$$\mathbf{W} = \text{diag} \left[w_1^2, w_2^2, \dots, w_D^2 \right] \quad (3).$$

[0057] Inputs that are judged to be close by the covariance function are likely to have similar outputs. A prediction is made by considering the covariance between the predictive case and all the training cases. For a noise-free input $\mathbf{x}^*$, the predictive distribution of $f(\mathbf{x}^*)$ is Gaussian with mean $\mu(\mathbf{x}^*)$ and variance $\sigma(\mathbf{x}^*)$:

$$\mu(\mathbf{x}^*) = \mathbf{k}(\mathbf{X}, \mathbf{x}^*)^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} \quad (4)$$

$$\sigma^2(\mathbf{x}^*) = k(\mathbf{x}^*, \mathbf{x}^*) - \mathbf{k}(\mathbf{X}, \mathbf{x}^*)^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}(\mathbf{X}, \mathbf{x}^*) \quad (5)$$

where $\mathbf{k}(\mathbf{X}, \mathbf{x}^*)$ is the $N \times 1$ vector of covariance functions between training inputs $\mathbf{X}$ and the new input $\mathbf{x}^*$. $\mathbf{K}$ is the $N \times N$ matrix of covariance functions between each pair of training inputs. σ_n^2 denotes the variance of Gaussian noise in training targets $\mathbf{y}$. σ_f , σ_n and $w_1, w_2 \dots w_D$ are hyperparameters to be trained in a Gaussian process.

[0058] FIG. 4 summarizes the procedures of using Gaussian processes for predictions in accordance with some embodiments. A Gaussian process is built upon training data, which can be, for example, sensor readings or metered data of a real system, or simulated data generated from complex models. As shown, the Gaussian process model takes new inputs and makes predictions with uncertainty. Depending of the type building systems to be monitored for fault detection, different features may be used to train a Gaussian process model that may be used in accordance with some embodiments. After the features are selected, the model may be trained to generate a baseline model that may be used for predication and fault detection. Examples of the feature selection, training, and prediction process in accordance with some embodiments are described in further detail below.

Feature Selection

[0059] The inventors have recognized and appreciated that some existing feature selection methods based on entropy, correlations coefficients, or mutual information may not be applicable to Gaussian process regression. Accordingly, in some embodiments, feature selection specifically targeted at Gaussian process regression may be used. For example feature selection in which hyperparameters that the Gaussian process regression learns may be used. In some embodiments, feature selection may include the following steps:

- Normalize all features and the target to $[-1, 1]$.
- Divide the dataset into training, validation and test sets. Use training and validation sets for feature selection. If the dataset is large and computational cost is a concern, select smaller training and validation sets for the purpose of feature selection.
- Include all the features to train a GP regression and derive the characteristic length-scale w_i that corresponds to each feature from the training set.
- Calculate the indicator of each feature as follows.

$$g(\mathbf{x}_i) = \frac{w_i}{\text{std}(\mathbf{x}_i)}$$

where $\mathbf{x}_i$ is the input vector of a certain feature, $i = 1, 2, \dots, d$ and d is the number of features. The indicator $g(\mathbf{x}_i)$ is defined as the characteristic length-scale of a feature divided by the standard deviation of all inputs on that feature dimension. The characteristic length-scale may determine how close two points have to be to influence each other. As $g(\mathbf{x}_i)$ is normalized by the variation of the training inputs on that feature dimension, it indicates how significant that feature is. If $g(\mathbf{x}_i)$ is small, it means small changes in input value of a certain feature will have a significant impact on output value. Therefore, a relatively small $g(\mathbf{x}_i)$ indicates that the corresponding feature is important in GP regression.

- Select $\mathbf{x}_{(1)}, \mathbf{x}_{(2)}, \dots, \mathbf{x}_{(d^*)}$ as follows:

$$\begin{aligned} \mathbf{x}_{(1)} &= \arg \min_X [g(\mathbf{x}_i)] \\ \mathbf{x}_{(2)} &= \arg \min_{X - \{\mathbf{x}_{(1)}\}} [g(\mathbf{x}_i)] \\ &\vdots \\ \mathbf{x}_{(d^*)} &= \arg \min_{X - \{\mathbf{x}_{(1)}, \mathbf{x}_{(2)}, \dots, \mathbf{x}_{(d^*-1)}\}} [g(\mathbf{x}_i)] \end{aligned}$$

[0060] In accordance with some embodiments, the feature with the smallest $g(\mathbf{x}_i)$ may be the first feature to select. The number of the features d^* may depend on training and validation accuracies, as well as computational cost. In general, if the training and/or validation accuracy is much lower than that when using all features, more features may be included in the model. Otherwise, fewer features may be used to further reduce computational time.

[0061] Some embodiments are directed to predicting whole building daily electricity, chilled water and steam consumption. In some implementations, the following features may be used in the Gaussian process model:

Electricity:

- Occupancy: A number between 0 and 1.0 indicates no occupants, 1 indicates normal occupancy. An estimate based on holidays, weekends and school academic calendar may be used.
- Day of week
- Day of year
- Week of year
- Heating degrees: if $15^{\circ}\text{C} - T_0 > 0$, then $= 15^{\circ}\text{C} - T_0$, else $= 0$. Assume that when outdoor temperature is above 15°C , no heating is needed.

Chilled water:

- Cooling degrees: if $T_0 - 12^{\circ}\text{C} > 0$, then $= T_0 - 12^{\circ}\text{C}$, else $= 0$. Assume that when outdoor temperature is below 12°C , no cooling is needed, which is true for many buildings. This is useful for daily prediction, because the average of hourly cooling degrees may be more preferable than average of hourly temperature.
- Dehumidification: if humidity ratio $- 0.00886 \text{ kg/kg} > 0$, then $= \text{humidity ratio} - 0.00886 \text{ kg/kg}$, else $= 0$. This feature may be particularly useful for daily chilled water prediction.
- Occupancy
- Day of week
- Solar radiation

Steam:

- Heating degrees
- Occupancy
- Day of week
- Day of year
- Week of year

The features described above are representative of a particular application in which whole building electricity, chilled water and steam consumption are monitored. However, it should be appreciated that different, including fewer or additional features may be selected depending on the particular building energy system monitoring application or user requirements.

[0062] FIG. 5 shows a portion of a user interface that enables a user to select features and training targets for use with some embodiments. As shown, a user may select features and targets by clicking on an interactive graph showing the nodes of the system. Once selection of the features and targets is finished, the model including the selected features may be trained.

[0063] FIG. 6 shows a portion of a user interface that enables a user to train the model using the selected features and training targets. A user may interact with the user interface to select a

training period. Once the training period is selected, users can click the “Train” button. After clicking the button, a window may pop up to ask users if validation needs to be performed. If yes, both training accuracy and validation accuracy may be displayed on the user interface. If validation is not needed, only training accuracy may be displayed. The accuracy displays may help users evaluate the model and chose the best model to be used for prediction.

[0064] FIG. 7 shows a portion of a user interface that enables a user to make predictions using the trained model. Once the model is trained, users can obtain the predication results using the Gaussian process algorithm with the measured and predicted values and the 95% confidence intervals. As shown, the accuracy of the prediction may also be displayed in the user interface.

Abnormality Threshold

[0065] From Equations 4 and 5 above, the predictive mean μ and the standard deviation σ of baseline energy consumption can be determined. Assuming that the system performs in the same way as it does during the time when the training data is collected, there is about 68% chance that the observed energy consumption falls within one standard deviation from the mean value. A threshold, $k\sigma$ can be defined. In some embodiments, a fault is detected when the observed energy consumption exceeds the baseline mean by $k\sigma$. The variable k may be configurable such that user can choose a different value for k based on a preference for fewer false positive errors (false alarms) or fewer false negative errors. In some embodiments, $k=2$. If $k < 2\sigma$, more false alarms may be observed, whereas a larger value for k may result in more faulty conditions being ignored.

Prediction Accuracy

[0066] The accuracy of baseline prediction depends on the quality of the training data, with the size of data sample being important. In some embodiments, one year of data is acquired and used as training data, as this amount of data will cover a typical range of weather and occupancy scenarios. Feature selection also influences the ability of the trained model to accurately predict faults in the system. In some embodiments, weather, occupancy and time features are included. Some examples of features that may be used in some embodiments, are described above. Feature selection also depends on the sample size and system type. For example, if the sample size does not cover the entire year, then day of year and weak of year might not be good features and extrapolation may be needed. If electricity consumption includes electricity used for chillers, then weather features may need to be added to the electricity prediction. If the

occupancy or function of the building has changed, or the building has been renovated, then the previous baseline model might no longer be applicable. In this situation, building managers may need to collect more data to train a new baseline model.

Component FDD

[0067] Some embodiments are directed to algorithms that combine physics-based principles and data-driven methods to detect faults in components of a building energy consumption system, such as an HVAC system. The algorithms take into consideration operation uncertainties of the building systems such as measurement uncertainties of sensor. As one example, a building FDD framework may monitor outdoor air (OA) control to detect and diagnose faults in the system. In this example, the FDD algorithm is configured to detect and diagnose the common OA control faults: OA damper stuck, economizer control fault, OA sensor fault, OA flow design fault, etc.

[0068] An important step for this OA control FDD algorithm is to compare the real OA ratio measurements with OA ratio setpoints. When the difference between the measured OA ratio and OA ratio setpoint exceeds a certain confidence threshold, a control fault is detected. In some embodiments, a 95% confidence region may be used.

[0069] FIG. 8 shows a plot of an OA ratio comparison between measured data and an OA ratio setpoint in accordance with some embodiments. As shown, an OA control fault is detected from 9:00 am to 13:00 pm on August 14th, 7:00 am to 13:00 pm on August 15th, and 7:00 am to 13:00 pm on August 16th.

[0070] In accordance with some embodiments, a fault diagnostics process may be executed to determine the fault sources. Any suitable fault diagnostics process may be used depending on the type of system or systems being monitored and embodiments are not limited in this respect. FIG. 9 shows an example of a fault diagnostics process for an OA damper stuck fault. The main reason for this fault is that the OA damper became stuck at a non-minimum position when the OA temperature was out of the economizer range.

[0071] In some embodiments configured to detect and diagnose component level OA control faults, the following features and data collection frequency may be used:

Feature selection:

- OA ratio: A number between 0 and 1.0 indicates no OA, 1 indicates all OA.

1. For the cases with direct OA and SA flow rate measurement, the OA ratio can be calculated directly from the OA flow rate and SA flow rate.

2. For the cases without OA and SA flow rate measurements, the OA temperature and SA temperature are needed to calculate the OA ratio, and the OA ratio can be estimated from:

$$r = \frac{T_m - T_r}{T_o - T_r},$$

where T_m is the mixed air temperature, T_r is the return air temperature and T_o is the outdoor air temperature. When the cooling coil is not active, the actual mixed air temperature is equal to the air handling unit supply air temperature minus the temperature rise due to fan heat ΔT_{fan} ,

$$T_m = T_s - \Delta T_{fan}.$$

- OA ratio setpoint: estimated from MA temperature setpoint or SA temperature setpoint:

$$r_{set} = \frac{T_{m,set} - T_r}{T_o - T_r},$$

where $T_{m,set}$ is the mixed air temperature setpoint.

The uncertainties in operations are considered as follows:

$$u(r) = \sqrt{\left(\frac{u^2(T_s) + u^2(\Delta T_{fan})}{(T_o - T_r)^2} \right)^2 + \left(\frac{u^2(T_r)(T_s - \Delta T_{fan} - T_o)^2}{(T_o - T_r)^4} \right)^2 + \left(\frac{u^2(T_o)(T_s - \Delta T_{fan} - T_r)^2}{(T_o - T_r)^4} \right)^2},$$

$$u(r_{set} - r) = \sqrt{\frac{u^2(T_s)}{(T_o - T_r)^2} + \left(\frac{(T_{s,set} - T_s)^2 (u^2(T_o) + u^2(T_r))}{(T_o - T_r)^4} \right)}$$

where $u(T_s)$, $u(T_o)$, and $u(T_r)$ are the measurement uncertainty for supply air temperature, outdoor air temperature and return air temperature and $u(\Delta T_{fan})$ is the estimation uncertainty of the temperature rise due to fan heat.

- Economizer control method:
 1. Temperature based control: high and low temperature limits for economize on/off
 2. Enthalpy based control: high and low enthalpy limits for economize on/off
- Cooling coil valve position

- OA damper position
- Minimum OA damper position
- Maximum OA damper position

Data collection frequency:

In some embodiments, a typical data collection frequencies of building automation system: 5 minutes, 15 minutes, and 1 hour, may be used.

Parametric Analysis

[0072] In accordance with some embodiments, parametric uncertainty is modeled since there is embedded uncertainty in occupancy estimation. In some cases, it is desired to investigate the impact of inputs on outputs by intentionally allowing inputs to vary in their domains. The impact of an HVAC control variable on energy consumption may be examined using GP regression. In one example, an office building having three air-cooled condensing units which supply refrigerant to three direct expansion air handling units (AHUs) was investigated. The terminal units were VAV boxes with reheat and its energy consumption was sub-metered. The electric energy consumption of three condensing units and the supply fans in three AHUs were metered individually at 15-min intervals. AHU sensor readings of supply air temperature and return air temperature were available from June to August 2012 at hourly intervals. The total building electric energy consumption consisted of both HVAC, lighting and plug load consumption. The gas consumption for heating and domestic hot water was also metered but not considered in this example, since only the summer months when AHU sensor readings were available were used.

[0073] The supply air temperature (SAT) set-point of an AHU is fixed. Different fixed set-point values were used for different seasons. As AHU SAT does not vary according to outdoor air temperature or zone thermal load, it is likely to be suboptimal. As shown in FIG. 10, AHU return air temperatures (RAT), which were close to zone average temperatures, vary according to outdoor air temperature (OAT). When OAT is low, zone air temperatures were far below the upper bound of indoor comfort level, which indicates excessive cooling, in particular for AHU1 and AHU2. AHU SAT may be further optimized. In building commissioning/retrofitting projects, it is worth knowing the cost-effectiveness beforehand. An estimate the impact of optimizing AHU SAT can assist decision-making. A physics-based simulation, e.g., using EnergyPlus, may be used to evaluate the importance of control variables and estimate energy

savings potential of optimization. However, as discussed above, physics-based simulations typically require detailed information of a building and its systems. Consequently, such simulations are labor-intensive and time-consuming. Consequently some embodiments are directed to a fast way to estimate the impact of a control variable on energy consumption. When historical data of system operations is available, some embodiments develop a data-driven model to replace the physics-based simulations in existing solutions to give an estimate of impact.

[0074] Table 2 below shows a distribution of measured AHU SATs and impact factor.

	Mean Value (°C/°F)	Standard Deviation (°C/°F)	Impact factor
AHU1	14.6 / 58.4	0.52 / 0.93	3.0%
AHU2	12.6 / 54.6	0.95 / 1.71	4.4%
AHU3	15.0 / 59.0	0.32 / 0.58	4.7%

[0075] Although the set-points are fixed, there is some variance in actual SAT control. Poor PID control, or insufficient or excessive cooling supply might account for the variance. The variation in actual SATs allows training a surrogate model based on historical data using GP regression. The Gaussian process model described above may be used to determine the distribution of energy consumption based on measured SATs. Then a variance in SAT(s) may be used to obtain new distributions. The additional variance in energy consumption caused by SAT variation provides an estimate of the impact of SAT on energy consumption.

[0076] Hourly data of outdoor air temperature, humidity, AHU SATs, RATs, sub-metered and total building electric energy consumption was available from June 12th to August 15th 2012. After removing unoccupied hours, there were 370 data points left. The first step was to use measured outdoor air temperature, humidity and AHU SATs as the features, and total building electric energy consumption as the target to train a Gaussian process regression model. In this example, the training R^2 was 0.96. The same set of data was then used to calculate the predictive distribution of each training target and determine the corresponding variance V_{y_i} of each point. Next a noise of $\sigma_{SAT} = 0.56$ °C (1°F) was added in the SAT of an AHU and the $V_{y_{noisy,i}}$ that accounts for noisy AHU SAT was determined. The impact factor was defined as the

sum of additional output variation caused by the input noise divided by the sum of predictive mean:

$$\text{Impact factor} = \frac{\sum_{i=1}^n (\sqrt{V_{y_{noisy},i}} - V_{y_i})}{\sum_{i=1}^n y_i}$$

[0077] The impact factor gives an estimate of the impact of the input variance on the output. A large value indicates a larger impact of the input. In this example, the impact factors of SATs of three AHUs are listed in Table 2. These factors give an estimate of the impact of a small change in each AHU SAT on total building electric energy consumption. SATs of AHU2 and AHU3 have larger impact than AHU1 SAT. This is consistent to sub-metered electric energy consumption of the AHUs and condensing units. As shown in FIG. 11, the electric energy consumption of AHU2 and AHU3 during the investigated time period was significantly larger than that of AHU1. Intuitively, the magnitude of impact factor of SAT is related to the energy consumption of that AHU, although not exactly proportional. Some embodiments that use Gaussian process regression provide a rapid estimate in order to assist early-decision making without the necessity of acquiring detailed information of a building and its system. The accuracy of the estimate depends on data quality, which directly affects the accuracy of surrogate model using Gaussian process regression.

Using Forecasting for “what-if” scenarios

[0078] Some embodiments are directed to performing what-if scenarios to enable a user to predict the energy cost saving potentials from fixing faults or optimizing system operation. That what-if scenarios use machine learning algorithms to develop forecasting models based on building operation data (real field measurements or simulation data) and uses these forecasting models to predict cost saving potentials. Based on the user’s selection of control upgrading variables or fault fixing scenarios, the “what-if” predicting algorithm calculates the energy cost saving potentials.

[0079] The cost saving potential data from simulations may be calculated from the difference between the original operation settings and the user desired settings. The cost savings can be whole year savings or savings for a certain time range. In some embodiments, the what-if predicting model uses two different approaches. A first approach bases energy balance on a theoretical calculation. A second approach uses machine learning forecasting. For the second approach, certain operation data may be collected and used to develop the forecasting models.

[0080] FIG. 12 shows an example of a user interface for a “what-if” model that may be used in accordance with some embodiments. As shown, the “what-if” model may be configured to predict the cost savings for fixing an OA damper stuck fault, supply air temperature setpoint optimization, supply air temperature reset, and building window upgrading.

Example system architecture

[0081] FIG. 13 schematically illustrates an example of a system architecture within with some embodiments may be used. The architecture includes building automation system (BAS) controllers connected to a BAS server. Each of the BAS server and one or more sensors connect to a smart building platform (SBP) via a building service interface (BSI). The SBP may also interface with third party services via one or more networks (e.g., the Internet).

[0082] FIG. 14 schematically illustrates the system architecture of FIG. 13 in which the sensor inputs and their associated BSI are collectively referred to as the system backend and the user input interface is referred to as the system frontend. In some embodiments, the system backend may be implemented as a web service model operating as a middleware system managing building resources (e.g., sensors, lights, actuators, etc.), collecting data, and interacting with other systems (e.g., control center and third party building control system). When the backend is implemented using a web service model, the building resources may be accessed via a URI using the HTTP protocol. Based on the backend web services, a web frontend may be used to present the service data to a user

[0083] FIG. 15 shows a block diagram of the system architecture of FIG. 13. The backend may be implemented as a web service (e.g., in Scala). The web service model is developed for cross platform compatibility to facilitate integration and development with other building automation/control systems. In some embodiments, a Representational State Transfer (REST) web service model is used to implement the backend providing a highly decoupled and scalable structure. The REST web services enable the frontend to interact with the SBP server and third parties to integrate the building services into their applications. The SBP REST web services provide access to resources via URI paths. To use a REST web service, a user computing device sends an HTTP request to a REST server and parses the response. In some embodiments, the response format is JavaScript Object Notation (JSON). Standard HTTP methods such as GET, PUT, POST and DELETE may be used. Because the REST web service is based on open standards, any web development language may be used to access the backend services.

[0084] The SBP is also configured to use one or more FDD algorithms, examples of which are described above, to perform fault detection and diagnostics using the system and sensor data received by the SBP.

[0085] In some embodiments, the frontend of the system is implemented as a web-based graphical user interface for users to interact with the system. Users can access the web interface using a web browser for data visualization, whole building consumption prediction, HVAC FDD, and what-if analysis. In one implementation, the visualizations may be implemented by the Data-Driven Documents library d3.js, which allows users to acquire more information through an interactive platform.

[0086] FIG. 16 illustrates an example of a front end dashboard that may be used in accordance with some embodiments. As shown, the dashboard displays some key performance metrics (e.g., cost, demand, and faults) at the top, the trends of the whole building consumption data in the middle, and a list of buildings in the system with their corresponding consumption data at the bottom. The whole building consumption data includes electricity, chilled water, and steam. For the trends of the whole building consumption data, the predictive mean and 95% confidence intervals is shown.

[0087] While various inventive embodiments have been described and illustrated herein, those of ordinary skill in the art will readily envision a variety of other means and/or structures for performing the function and/or obtaining the results and/or one or more of the advantages described herein, and each of such variations and/or modifications is deemed to be within the scope of the inventive embodiments described herein. More generally, those skilled in the art will readily appreciate that all parameters, dimensions, materials, and configurations described herein are meant to be exemplary and that the actual parameters, dimensions, materials, and/or configurations will depend upon the specific application or applications for which the inventive teachings is/are used. Those skilled in the art will recognize, or be able to ascertain using no more than routine experimentation, many equivalents to the specific inventive embodiments described herein. It is, therefore, to be understood that the foregoing embodiments are presented by way of example only and that, equivalents thereto, inventive embodiments may be practiced otherwise than as specifically described. Inventive embodiments of the present disclosure are directed to each individual feature, system, article, material, kit, and/or method described herein. In addition, any combination of two or more such features, systems, articles, materials, kits,

and/or methods, if such features, systems, articles, materials, kits, and/or methods are not mutually inconsistent, is included within the inventive scope of the present disclosure.

[0088] Also, the technology described herein may be embodied as a method, of which an example has been provided. The acts performed as part of the method may be ordered in any suitable way. Accordingly, embodiments may be constructed in which acts are performed in an order different than illustrated, which may include performing some acts simultaneously, even though shown as sequential acts in illustrative embodiments.

[0089] All definitions, as defined and used herein, should be understood to control over dictionary definitions, definitions in documents incorporated by reference, and/or ordinary meanings of the defined terms.

[0090] The indefinite articles “a” and “an,” as used herein, unless clearly indicated to the contrary, should be understood to mean “at least one.”

[0091] The phrase “and/or,” as used herein, should be understood to mean “either or both” of the elements so conjoined, i.e., elements that are conjunctively present in some cases and disjunctively present in other cases. Multiple elements listed with “and/or” should be construed in the same fashion, i.e., “one or more” of the elements so conjoined. Other elements may optionally be present other than the elements specifically identified by the “and/or” clause, whether related or unrelated to those elements specifically identified. Thus, as a non-limiting example, a reference to “A and/or B”, when used in conjunction with open-ended language such as “comprising” can refer, in one embodiment, to A only (optionally including elements other than B); in another embodiment, to B only (optionally including elements other than A); in yet another embodiment, to both A and B (optionally including other elements); etc.

[0092] As used herein, “or” should be understood to have the same meaning as “and/or” as defined above. For example, when separating items in a list, “or” or “and/or” shall be interpreted as being inclusive, i.e., the inclusion of at least one, but also including more than one, of a number or list of elements, and, optionally, additional unlisted items. In general, the term “or” as used herein shall only be interpreted as indicating exclusive alternatives (i.e. “one or the other but not both”) when preceded by terms of exclusivity, such as “either,” “one of,” “only one of,” or “exactly one of.”

[0093] As used herein, the phrase “at least one,” in reference to a list of one or more elements, should be understood to mean at least one element selected from any one or more of the elements in the list of elements, but not necessarily including at least one of each and every

element specifically listed within the list of elements and not excluding any combinations of elements in the list of elements. This definition also allows that elements may optionally be present other than the elements specifically identified within the list of elements to which the phrase “at least one” refers, whether related or unrelated to those elements specifically identified. Thus, as a non-limiting example, “at least one of A and B” (or, equivalently, “at least one of A or B,” or, equivalently “at least one of A and/or B”) can refer, in one embodiment, to at least one, optionally including more than one, A, with no B present (and optionally including elements other than B); in another embodiment, to at least one, optionally including more than one, B, with no A present (and optionally including elements other than A); in yet another embodiment, to at least one, optionally including more than one, A, and at least one, optionally including more than one, B (and optionally including other elements); etc.

[0094] What is claimed is:

CLAIMS

1. A computer-implemented method of predicting resource usage within a building, the method comprising:

receiving data values representing usage measurements of one or more resources, the data values being produced by one or more sensors located within and/or near the building;

identifying a plurality of features, each feature of the plurality of features being a measureable property of one of the one or more resources and thereby associated with the usage measurements of the one of the one or more resources;

training, using at least one processor, a first Gaussian process regression model on a first subset of the data values by optimizing first hyperparameters of the first Gaussian process regression model, wherein each of the first hyperparameters is associated with one of the plurality of features;

calculating, using the at least one processor, based on the first Gaussian process regression model, an indicator value for each feature of the plurality of features, the indicator value indicating an extent to which variation in values of the feature correlates to variation in the usage measurements produced by the one or more sensors and associated with the feature;

training, using the at least one processor, a second Gaussian process regression model on a second subset of the data values by optimizing second hyperparameters of the second Gaussian process regression model, wherein each of the second hyperparameters is associated with one of a subset of the plurality of features, wherein features of the subset are selected from amongst the plurality of features based, at least in part, on said indicator values calculated for the features of the subset; and

predicting, using the at least one processor, usage of a first resource of the one or more resources based on the trained second Gaussian process regression model.

2. The method of claim 1, wherein predicting the usage of the first resource comprises predicting a usage amount of the first resource and an estimate of uncertainty associated with the predicted usage amount of the first resource.

3. The method of claim 1, further comprising detecting an anomaly by comparing additional usage measurements of the first resource with the predicted usage of the first resource.

4. The method of claim 3, wherein detecting the anomaly is further based on an estimate of uncertainty associated with the predicted usage of the first resource.
5. The method of claim 4, wherein detecting the anomaly is based on identifying that the additional usage measurements of the first resource exceed a threshold confidence interval defined by the estimate of uncertainty associated with the predicted usage of the first resource.
6. The method of claim 1, wherein the one or more resources includes one or more of electricity, water, steam, oil and/or gas.
7. The method of claim 1, wherein the plurality of features include a first feature that is a measureable property of the first resource, and wherein the first feature indicates usage of the first resource by one of: day of the year, week of the year, or day of the week.
8. The method of claim 1, wherein the first hyperparameters are characteristic length scales in covariance functions of the first Gaussian process regression model each associated with the one of the plurality of features.
9. The method of claim 1, further comprising predicting, using the at least one processor, usage of a second resource of the one or more resources based on the trained second Gaussian process regression model.
10. The method of claim 1, further comprising displaying the predicted usage amount of the first resource to a user via at least one display.
11. The method of claim 2, further comprising calculating, using the at least one processor, an impact factor of the first resource based at least in part on the predicted usage amount of the first resource and the estimate of uncertainty associated with the predicted usage amount of the first resource.

12. The method of claim 11, wherein calculating the impact factor of utilization of the first resource is further based on an amount of noise added to the estimate of uncertainty associated with the predicted usage amount of the first resource.

13. At least one computer-readable medium comprising instructions that, when executed by at least one processor, perform a method of predicting resource usage within a building, the method comprising:

receiving data values representing usage measurements of one or more resources, the data values being produced by one or more sensors located within and/or near the building;

identifying a plurality of features, each feature of the plurality of features being a measureable property of one of the one or more resources and thereby associated with the usage measurements of the one of the one or more resources;

training, using the at least one processor, a first Gaussian process regression model on a first subset of the data values by optimizing first hyperparameters of the first Gaussian process regression model, wherein each of the first hyperparameters is associated with one of the plurality of features;

calculating, using the at least one processor, based on the first Gaussian process regression model, an indicator value for each feature of the plurality of features, the indicator value indicating an extent to which variation in values of the feature correlates to variation in the usage measurements produced by the one or more sensors and associated with the feature;

training, using the at least one processor, a second Gaussian process regression model on a second subset of the data values by optimizing second hyperparameters of the second Gaussian process regression model, wherein each of the second hyperparameters is associated with one of a subset of the plurality of features, wherein features of the subset are selected from amongst the plurality of features based, at least in part, on said indicator values calculated for the features of the subset; and

predicting, using the at least one processor, usage of a first resource of the one or more resources based on the trained second Gaussian process regression model.

14. The at least one computer-readable medium of claim 13, wherein predicting the usage of the first resource comprises predicting a usage amount of the first resource and an estimate of uncertainty associated with the predicted usage amount of the first resource.

15. The at least one computer-readable medium of claim 13, wherein the method further comprises detecting an anomaly by comparing additional usage measurements of the first resource with the predicted usage of the first resource.
16. The at least one computer-readable medium of claim 13, wherein the one or more resources includes one or more of electricity, water, steam, oil and/or gas.
17. The at least one computer-readable medium of claim 13, wherein the first hyperparameters are characteristic length scales in covariance functions of the first Gaussian process regression model each associated with the one of the plurality of features.
18. A system, comprising:
at least one processor; and
at least one computer-readable medium comprising instructions that, when executed by the at least one processor, perform a method of predicting resource usage within a building, the method comprising:
receiving data values representing usage measurements of one or more resources, the data values being produced by one or more sensors located within and/or near the building;
identifying a plurality of features, each feature of the plurality of features being a measureable property of one of the one or more resources and thereby associated with the usage measurements of the one of the one or more resources;
training, using the at least one processor, a first Gaussian process regression model on a first subset of the data values by optimizing first hyperparameters of the first Gaussian process regression model, wherein each of the first hyperparameters is associated with one of the plurality of features;
calculating, using the at least one processor, based on the first Gaussian process regression model, an indicator value for each feature of the plurality of features, the indicator value indicating an extent to which variation in values of the feature correlates to variation in the usage measurements produced by the one or more sensors and associated with the feature;
training, using the at least one processor, a second Gaussian process regression model on a second subset of the data values by optimizing second hyperparameters of

the second Gaussian process regression model, wherein each of the second hyperparameters is associated with one of a subset of the plurality of features, wherein features of the subset are selected from amongst the plurality of features based, at least in part, on said indicator values calculated for the features of the subset; and
predicting, using the at least one processor, usage of a first resource of the one or more resources based on the trained second Gaussian process regression model.

19. The system of claim 18, wherein predicting the usage of the first resource comprises predicting a usage amount of the first resource and an estimate of uncertainty associated with the predicted usage amount of the first resource.

20. The system of claim 18, wherein the method further comprises detecting an anomaly by comparing additional usage measurements of the first resource with the predicted usage of the first resource.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG 6




FIG. 6 (Cont.)




FIG. 7




FIG. 7(Cont.)





FIG. 8

**FIG. 9**





FIG. 10

**FIG. 11**




FIG. 12




FIG. 13




FIG. 14




FIG. 15

Harvard CGBC Smart Building Platform
Gund Hall ΔWendo
05/01/16 - 04/15/16

Dashboard

\$ 1310 Cost

20%

View Details

157 kWh Consumed

5%

View Details

Trend

Electricity (kWh)

Predicted

Measured

91% Confidence region

Buildings

Show 10 entries

Name	Cost	Electricity (kWh)	Chilled Water	Gas (Btu)	Steam (Use)
18-18A Median Street	563	228	141	1955	901
1721 Cambridge Street	248	79	216	3546	408
Lithauer Center	1225	208	929	168	248
Biological Laboratory	2362	49	572	236	152
7 Grant Street	1575	274	144	1427	650
University Hall	1054	206	749	1904	630
Perkins Hall	1115	18	406	2729	758
20 Garden Street	4909	196	417	1664	27
54 Dunsler Street	1291	131	564	1051	810
Standish Hall Winthrop House	1091	151	432	2353	300

Showing 1 to 10 of 151 entries

Previous 1 2 3 4 5 16 Next

FIG. 16

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US 17/66687

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06Q 50/06 (2018.01)

CPC - G06Q 50/06, G06Q 10/06315, G06F 17/30598, G06Q 30/0202, G06Q 10/00, G06Q 10/063, G06Q 10/0631, G06F 11/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History Document

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

See Search History Document

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

See Search History Document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2015/0317589 A1 (The Trustees Of Columbia University In The City Of New York) 05 November 2015 (05.11.2015) entire document especially Abstract, para [0007]-[0008], para [0101], para [0108]-[0110], para [0119]-[0124], para [0140]-[0143], para [0147]-[0149], para [0180]-[0182]	1-20
Y	US 2010/0174514 A1 (Melkumyan et al.) 08 July 2010 (08.07.2010) entire document especially Abstract, para [0008]-[0012], para [0061]-[0064]	1-20
A	US 2016/0146493 A1 (International Business Machines Corporation) 26 May 2016 (26.05.2016) entire document	1-20

☐ Further documents are listed in the continuation of Box C.

☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

09 February 2018

Date of mailing of the international search report

07 MAR 2018

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774