

Estimation of building occupancy levels through environmental signals deconvolution

Afroz Ebadat, Giulio Bottegal, Damiano Varagnolo, Bo Wahlberg, Karl H. Johansson
ACCESS Linnaeus Centre, School of Electrical Engineering, KTH Royal Institute of Technology
{ ebadat | bottegal | damiano | bo | kallej }@kth.se

Abstract

We address the problem of estimating the occupancy levels in rooms using the information available in standard HVAC systems. Instead of employing dedicated devices, we exploit the significant statistical correlations between the occupancy levels and the CO₂ concentration, room temperature, and ventilation actuation signals in order to identify a dynamic model. The building occupancy estimation problem is formulated as a regularized deconvolution problem, where the estimated occupancy is the input that, when injected into the identified model, best explains the currently measured CO₂ levels. Since occupancy levels are piecewise constant, the zero norm of occupancy is plugged into the cost function to penalize non-piecewise constant inputs. The problem then is seen as a particular case of *fused-lasso* estimator by relaxing the zero norm into the ℓ_1 norm. We propose both online and offline estimators; the latter is shown to perform favorably compared to other data-based building occupancy estimators. Results on a real testbed show that the MSE of the proposed scheme, trained on a one-week-long dataset, is half the MSE of equivalent Neural Network (NN) or Support Vector Machine (SVM) estimation strategies.

1 Introduction

In recent years we have witnessed an emerging interest in improving buildings energy efficiency by optimizing the performance of Heating, Ventilation and Air Conditioning (HVAC) systems. A common opinion is that a key enabling factor to achieve this goal is to have reliable instruments for detecting the occupancy levels in rooms [1, 2, 3]. In fact, the presence of occupants has a direct effect on the

air quality indexes (temperature, CO₂ and humidity levels); moreover, conditioning unused spaces usually translates into energy waste.

Clearly, the simplest way to obtain information regarding room occupancy is to employ dedicated hardware, such as cameras and RFID tags. However, this usually requires additional costs, poses privacy concerns, and might be infeasible in some old buildings. Furthermore, authors' direct experience indicates that the reliability of some standard off-the-shelf devices may be insufficient for the employment in HVAC control systems.

The natural question arising is what information on room building occupancy can be successfully gathered using indirectly the measurement devices usually present in standard HVAC systems. In other words, it is interesting to understand whether statistical processing of environmental signals such as CO₂ concentration, temperature and actuation signals may lead to occupancy estimators that are sufficiently valid to substitute hardware-based people counters. In this paper we precisely address this problem, and focus on how to accurately estimate the number of occupants in a room by processing CO₂, temperature and HVAC actuation levels.

Literature review

There exists a quite rich literature on methods and devices for the inference of occupancy levels in rooms and buildings. A first categorization divides it in two main branches: *hardware-based* approaches and *model-based* approaches.

The first category comprises strategies based on dedicated hardware such as cameras, RFIDs, etc. [4, 5, 6, 7, 8, 9, 10, 11, 12]. As mentioned before, these methods present some drawbacks that might make them non-suitable in certain scenarios. Instead, the second category of methods adopts procedures that indirectly infer occupancy levels using suitable models of the dynamics for some available environmental signals. These techniques typically exploit models that relate the occupancy patterns to the measured environmental signals; the occupancy levels are then estimated by applying proper inverse mappings starting from the available measurements. This category of methods can be furthermore divided in two classes: *physics-based* methods, where the model is derived exploiting knowledge of the underlying physical laws, and *identification-based* methods, where the model is derived using data-driven techniques.

As for the physics-based procedures, the rationale is usually to connect number of occupants, CO₂ concentration,

This work is supported by the European Institute of Technology (EIT) Information and Communication Technology (ICT) Labs, the Swedish Energy Agency, the Swedish Governmental Agency for Innovation Systems (VINNOVA) and the Knut and Alice Wallenberg Foundation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

BuildSys'13, November 14-15, 2013, Rome, Italy
Copyright © 2013 ACM TODO ...\$10.00

temperature and humidity with mass balance equations or first principles considerations [13, 14, 15, 16]. Opposed to the previous methods, identification-based approaches construct models from data-sets of past measured data. Successful approaches exploit machine learning techniques such as Support Vector Machines (SVMs), Neural Networks (NNs) and Hidden Markov Models (HMMs), using CO₂-based features (i.e., averages of the signals in time, first and second order temporal differences) [17, 18].

Statement of contributions

We propose a novel strategy that, with respect to the literature analyzed before, fall into the family of the identification-based approaches. More precisely, we identify Linear Time Invariant (LTI) dynamic models exploiting Prediction Error Method (PEM) approaches. Notice that identification of systems under feedback is not an easy task. In particular, we analyze the performance of two different identification algorithms, relying on parametric and nonparametric techniques respectively. Once the model is obtained, the building occupancy estimation is formulated as a deconvolution problem, i.e., an inverse problem where the input of a system is estimated from measured output data. In particular, we leverage the fact that the occupancy signal is *piecewise constant* and integer. This estimation problem is formulated in both an online and an offline version; while the first one has evident application in HVAC control systems, the second one can be exploited for deducing occupancy pattern models in large buildings.

The proposed scheme assumes the availability of measurements of temperature, CO₂ concentration and HVAC actuation levels at any time. This appears reasonable since this is exactly the information processed by standard Indoor Air Quality (IAQ) management systems. Indeed, these systems generally measure temperature and CO₂ levels to decide the heating, venting and cooling actuation levels. Furthermore, to identify the LTI model, we assume knowledge of the true occupancy levels for a short and well-defined period of time. We remark that no a-priori knowledge on the characteristics of the room is assumed.

As part of our work, we assess the importance of considering the HVAC actuation levels for occupancy estimation purposes. More precisely, we numerically evaluate the variation of the performance of the estimators when the amount of fresh air inlet in the room is neglected. As expected, exploitation of this information is beneficial for the quality of the estimation process. We eventually compare our approach with SVM- and NN-based algorithms, to understand to what extent the novel procedures improve the current state of the art. The approach is evaluated on real testbed and outcomes show that the proposed approach can reach, after a training period of one week, an estimation accuracy of ~88%.


Structure of the manuscript

Section 2 defines the mathematical problem under consideration, and the methodology used to solve it. Sections 3 and 4 describe the two main parts of the novel estimators, namely how to build the model of the room from a training set, and how to exploit the model for estimation purposes. Section 5 reports the SVM- and NN-based schemes used for comparison purposes. Section 6 describes the performance

indexes considered, the experimental setup and the results of the estimation processes. Section 7 then wraps our conclusions, remarks, and ideas for future directions of development.

2 Problem definition and methodology

We consider the following schematic representation of the dynamic behavior of the concentration of the CO₂ and temperature in a room under well-mixed air assumptions, i.e., in a room where these quantities are assumed to be spatially constant.



In the above scheme $c(k)$ represents the concentration of CO₂, $t(k)$ the temperature, $v(k)$ the amount of injected fresh air (venting) and $o(k)$ the occupancy, every one at time k . $G(s)$ represents a linear time-invariant system, i.e., a dynamic map transforming disturbances, ventilation and building occupancy levels into temperature and CO₂ concentration signals. We assume to be able to collect samples of these quantities (except for the disturbances) at specific time instants, and no knowledge of the model $G(s)$.

We are then interested in defining an effective model-based estimator of $o(k)$ which exploits the measurements $c(k), c(k-1), \dots, t(k), t(k-1), \dots, v(k), v(k-1), \dots$. To this end, we define two different phases.

Training phase. Here we assume to be able to measure T_r samples of $c(k), t(k), v(k)$, and, thanks to a temporary people counter, also $o(k)$. With this information we build an estimate $\hat{G}(s)$ of $G(s)$ (see next Section).

Test phase (either online or offline). Here we assume to be able to measure only $c(k), t(k)$, and $v(k)$. With this information, and exploiting the estimated model of the room $\hat{G}(s)$, we build an estimate $\hat{o}(k)$ of $o(k)$. When the estimation problem is solved on a fixed dataset of size T_s , we shall say that we are employing an offline estimator. On the other hand, an online estimator is defined when $\hat{o}(k)$ is estimated at each time instant.

The first problem is said to be a *system identification* problem, while the second is usually called a *deconvolution* problem. In the following sections we introduce our approaches for solving these two problems.

3 Identification of the model of the room

Assume we are in training phase, and thus to aim at building an estimate $\hat{G}(s)$ of $G(s)$ starting from a dataset of measurements of $c(k), t(k), v(k)$, and $o(k)$.

As in [19, 20, 21, 22, 23], we assume the environmental signals to be stationary and the dynamics of the room to be LTI, and thus of the form

$$\begin{bmatrix} c(k) \\ t(k) \end{bmatrix} = \begin{bmatrix} \mathbf{H}_c(q^{-1}) \\ \mathbf{H}_t(q^{-1}) \end{bmatrix} \begin{bmatrix} c(k-1) \\ t(k-1) \\ v(k-1) \\ o(k-1) \end{bmatrix} + \begin{bmatrix} w_c(k) \\ w_t(k) \end{bmatrix} \quad (1)$$

where $w^c(k), w^t(k)$ are white Gaussian noises and

$$\begin{aligned} \mathbf{H}_c(q^{-1}) &:= [\mathbf{H}_c^c(q^{-1}) \ \mathbf{H}_c^t(q^{-1}) \ \mathbf{H}_c^v(q^{-1}) \ \mathbf{H}_c^o(q^{-1})], \\ \mathbf{H}_t(q^{-1}) &:= [\mathbf{H}_t^c(q^{-1}) \ \mathbf{H}_t^t(q^{-1}) \ \mathbf{H}_t^v(q^{-1}) \ \mathbf{H}_t^o(q^{-1})] \end{aligned}$$

are matrix transfer functions where, without loss of generality, all the entries are assumed to be polynomials of the same order.

To estimate the transfer functions $\mathbf{H}_c(q^{-1})$ and $\mathbf{H}_t(q^{-1})$, that fully represent the (discretized) model of the system, we consider a classical PEM paradigm, i.e., we consider a one-step-ahead predictor of the form

$$\begin{bmatrix} \hat{c}(k|k-1) \\ \hat{t}(k|k-1) \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{H}}_c(q^{-1}) \\ \hat{\mathbf{H}}_t(q^{-1}) \end{bmatrix} \begin{bmatrix} c(k-1) \\ t(k-1) \\ v(k-1) \\ o(k-1) \end{bmatrix} \quad (2)$$

and then obtain $\hat{\mathbf{H}}_c(q^{-1})$ and $\hat{\mathbf{H}}_t(q^{-1})$ as those transfer functions that minimize the variance of the prediction errors $c(k) - \hat{c}(k|k-1)$ and $t(k) - \hat{t}(k|k-1)$ on the data collected during the training phase. From (2) it follows that the predictors $\hat{c}(k|k-1)$ and $\hat{t}(k|k-1)$ exploit the same information of the past; moreover, statistical analysis reveals that the signal mostly correlated with the occupancy is the CO₂ level. Thus, in the following we will consider only $\hat{c}(k|k-1)$ and thus focus on the identification of $\hat{\mathbf{H}}_c(q^{-1})$. To this end we consider the two following alternative PEM-based identification approaches:

parametric identification (P). Here the model is determined by parameters entering in $\mathbf{H}_c(q^{-1})$ and $\mathbf{H}_t(q^{-1})$ within a certain model structure, determined either by using model selection criteria such as AIC and BIC [24], or by adopting some physics-based considerations. As suggested in [14, 23], here we consider the latter strategy and exploit Autoregressive exogenous (ARX) models [24], where H_c^v , H_c^o , H_c^t and H_c^c are polynomials of order 1;

nonparametric identification (NP). Here, instead of imposing the order of the polynomials H_c^v , H_c^o , H_c^t and H_c^c and then searching for their optimal coefficients, we identify the corresponding impulse responses that best describe the evolution of the system during the training phase. Then, the coefficients of these impulse response (truncated to a fixed large number p) are used to form the aforementioned polynomials (see [25, 26] for more details).

Independently of the considered PEM method used, with the proposed identification step we obtain a predictor of the form

$$\hat{c}(k|k-1) = \hat{\mathbf{H}}_c(q^{-1}) \begin{bmatrix} c(k-1) \\ t(k-1) \\ v(k-1) \\ o(k-1) \end{bmatrix} \quad (3)$$

where

$$\hat{\mathbf{H}}_c(q^{-1}) = \begin{bmatrix} \hat{H}_c^c(q^{-1}) & \hat{H}_c^t(q^{-1}) & \hat{H}_c^v(q^{-1}) & \hat{H}_c^o(q^{-1}) \end{bmatrix} \quad (4)$$

is the matrix of the identified polynomials.

4 Deconvolution of the occupancy levels

Assume that we are in the test phase, where direct measurements of the building occupancy level $o(k)$ are no more available. Below we describe how, given estimates of the room dynamics model \hat{H}_c^c , \hat{H}_c^t , \hat{H}_c^v , \hat{H}_c^o obtained before, it is possible to build a function of the measurements $c(k)$, $v(k)$, and $t(k)$ that works as an estimator $\hat{o}(k)$ of $o(k)$.

Let then

$$e(k) := c(k) - \hat{c}(k|k-1) \quad (5)$$

be the prediction error of the CO₂ level, Gaussian by hypothesis [27]. Substituting (3) into (5) and rearranging properly one finds

$$\begin{aligned} \hat{H}_c^o(z^{-1}) o(k-1) &= \\ &= c(k) - \begin{bmatrix} \hat{H}_c^c(z^{-1}) & \hat{H}_c^t(z^{-1}) & \hat{H}_c^v(z^{-1}) \end{bmatrix} \begin{bmatrix} c(k-1) \\ t(k-1) \\ v(k-1) \end{bmatrix} - e(k), \end{aligned} \quad (6)$$

an identity where the unknowns are only $o(k)$ and $e(k)$, since

$$\tilde{c}(k) := c(k) - \begin{bmatrix} \hat{H}_c^c(z^{-1}) & \hat{H}_c^t(z^{-1}) & \hat{H}_c^v(z^{-1}) \end{bmatrix} \begin{bmatrix} c(k-1) \\ t(k-1) \\ v(k-1) \end{bmatrix}$$

can be computed given the available information.

With this new definition (6) can be rewritten as

$$\tilde{c}(k) = \hat{H}_c^o(q^{-1}) o(k-1) + e(k). \quad (7)$$

The above equation reformulates the problem of estimating the unknown $o(\cdot)$ as a deconvolution problem, i.e., of estimating the unknown occupancy signal as that input that best describes the observed output, given the knowledge of the transfer function of the system [28]. Since $e(k)$ is assumed to be white and Gaussian, the natural approach to this problem would be to employ a Least-Squares (LS) estimator of $o(\cdot)$, since this would minimize the overall variance of the estimation error [24, Chap. 7]. However, the estimates obtained in this way are usually unsatisfactory, since they do not really reflect a suitable building occupancy pattern. For this reason, we consider a cost function for the deconvolution problem that takes into account the prior information on the building occupancy signal, that is:

- $o(k)$ is a non-negative integer signal;

- $o(k)$ is piecewise constant.

We thus formulate the deconvolution problem as “find that least-changing positive piecewise constant input signal that minimizes the mismatch between the estimated and measured outputs of the system”. As underlined by the structure of this section, this problem can be solved both offline and online.

4.1 Offline estimation of the building occupancy levels

Starting from the previous findings and assumptions, the offline estimation problem, which is motivated by situations where one wants to construct models of building occupancy flows in buildings, can be formulated as follows.

Let the polynomial $\hat{H}_c^o(q^{-1})$ be $\hat{H}_c^o(q^{-1}) = h_0 + h_1 q^{-1} + \dots + h_p q^{-p}$, where $p = 1$ when ARX model is adopted, and

let the test set be indexed by the time instants $0, \dots, T_{ts}$. Consider the auxiliary notation

$$\hat{H} := \begin{bmatrix} h_p & \cdots & h_0 & 0 \\ & \ddots & & \\ 0 & & h_p & \cdots & h_0 \end{bmatrix} \in \mathbb{R}^{T_{ts}-p \times T_{ts}}$$

$$\mathbf{o} := \begin{bmatrix} o(0) \\ \vdots \\ o(T_{ts}-1) \end{bmatrix}, \quad \tilde{\mathbf{c}} := \begin{bmatrix} \tilde{c}(p+1) \\ \vdots \\ \tilde{c}(T_{ts}) \end{bmatrix}$$

$$\Delta o(i) := o(i) - o(i-1), \quad \Delta \mathbf{o} := [\Delta o(1), \dots, \Delta o(T_{ts}-1)].$$

The offline estimation problem can then be expressed mathematically as

$$\hat{\mathbf{o}} = \arg \min_{\mathbf{o} \in \mathbb{N}_+^{T_{ts}}} \|\tilde{\mathbf{c}} - \hat{H} \mathbf{o}\|_2^2 + \lambda \|\Delta \tilde{\mathbf{o}}\|_0, \quad (8)$$

where:

- $\hat{\mathbf{o}}$ is a T_{ts} -dimensional vector with the estimated values of building occupancy at the time instants $0, \dots, T_{ts}-1$;
- the first summand on the RHS represents the LS estimator of the building occupancy, that tries to match the estimated and measured outputs of the system;
- $\|\cdot\|_0$, the zero norm, counts the number of variations of the candidate inputs, thus penalizing non-piecewise constant candidate inputs;
- λ is a regularization parameter that trades off the two previous terms and that is discussed in details in Section 4.3.

Unfortunately problem (8) is a non-convex non-linear integer program, and cannot be solved efficiently even by the most advanced numerical optimization procedures. To circumvent this computational drawback, we pose two relaxations. First, we substitute the zero norm with the ℓ_1 -norm [29, Sec. 3.4]; then, we extend the domain of the plausible inputs to $\mathbb{R}_+^{T_{ts}}$ instead of $\mathbb{N}_+^{T_{ts}}$. Hence, we formulate the problem of estimating the building occupancy levels as

$$\hat{\mathbf{o}} = \left[\arg \min_{\mathbf{o} \in \mathbb{R}_+^{T_{ts}}} \|\tilde{\mathbf{c}} - \hat{H} \mathbf{o}\|_2^2 + \lambda \|\Delta \tilde{\mathbf{o}}\|_1 \right], \quad (9)$$

where the $[\cdot]$ denotes the vector-wise rounding operator. Problem (9) can also be seen as a particular case of *fused-lasso* estimator [30], where the solution is searched among *sparse and smooth* regressor vectors.

4.2 Online estimation of the building occupancy levels

As mentioned before, the online estimation problem is motivated by situations where current building occupancy levels are used for actively controlling HVAC systems, e.g., [31]. It can be derived from the previously introduced offline estimator (9) by introducing some modifications. In particular, instead of using a fixed dataset (the test set), at each time instant we employ N data samples of each signals, from $k-N+1$ to k , with $N > p$. Then, by introducing

$\hat{H} \in \mathbb{R}^{N-p \times N}$ similarly as before, and

$$\mathbf{o} := \begin{bmatrix} o(0) \\ \vdots \\ o(N-1) \end{bmatrix}, \quad \tilde{\mathbf{c}} := \begin{bmatrix} \tilde{c}(p+1) \\ \vdots \\ \tilde{c}(N) \end{bmatrix}$$

$$\Delta o(i) := o(i) - o(i-1), \quad \Delta \mathbf{o} := [\Delta o(1), \dots, \Delta o(N-1)],$$

the estimate of the room occupancy at time $k-1$ is given by

$$\hat{\mathbf{o}}(k-1) = \left[\arg \min_{\mathbf{o} \in \mathbb{R}_+^N} \|\tilde{\mathbf{c}} - \hat{H} \mathbf{o}\|_2^2 + \lambda \|\Delta \tilde{\mathbf{o}}\|_1 \right]. \quad (10)$$

In (10), $\hat{\mathbf{o}}(k-1)$ is an N -dimensional vector with the estimated values of occupancy at the time instants $k-1, \dots, k-N$; one can consider only its first entry, i.e., $\hat{o}(k-1)$.

The parameter N plays an important role in (10), since it defines the amount of data employed for estimating $\hat{\mathbf{o}}(k-1)$ (and in particular $\hat{o}(k-1)$) at each time instant. Clearly, a large value of N yields more accurate estimates, since more information is used. In fact, note that the ratio between the number of data samples and the number of estimated values, given by $\frac{N-p}{N}$, tends to 1 as N grows large. However, a large value of N brings computational issues which could make the computation of (10) too slow for online operations. Thus, a good choice of N should be made by considering both these aspects; this is discussed in details in Section 6.3.

4.3 Finding the optimal regularization parameter λ

The solutions to Problems (9) and (10) depend on the choice of the regularization parameter λ . This weighting factor is indeed a design parameter that plays a crucial role in the estimation process, as it tunes the typical variability of the room occupancy signal. In other words, large values of λ penalize changes in the value of estimated occupancy, leading to estimates that are constant for long periods of times. On the other hand, small values of λ lead to occupancy estimates that vary frequently and behave similarly to the outcomes of the LS estimator (which is obtained by setting $\lambda = 0$).

A reasonable choice of λ is given by the value of such parameter that gives the best estimation performance during the training phase. In other words, the estimation of λ is computed first by plugging into (9) the measurements of $c(k)$, $i(k)$ and $v(k)$ collected *during the training phase*. Then, letting $\hat{\mathbf{o}}(\lambda)$ be solution of (9) applied to the training set with that particular λ , we compute the optimal regularization parameter as

$$\hat{\lambda} = \arg \min_{\lambda \in \mathbb{R}_+} \|\hat{\mathbf{o}}(\lambda) - \mathbf{o}\|_2^2, \quad (11)$$

where \mathbf{o} is the vector of the occupancy levels measured during the training phase.

5 Alternative methods

The algorithm proposed above is based on the assumption of having LTI models. If such hypothesis is neglected, some alternative methods can be utilized. Here we consider two of them for comparison purposes.

5.1 Estimation of building occupancy levels using Support Vector Machine (SVM)

In their basic form, SVMs perform classification tasks as follows: given a dataset \mathcal{D} of samples (\mathbf{x}_k, y_k) for $k = 0, \dots, N$ with $\mathbf{x}_k \in \mathbb{R}^n$ and $y_k \in \{-1, +1\}$, try to find a hyperplane in \mathbb{R}^{n+1} that: (i) separates the points of the form $(\mathbf{x}_k, +1)$ from those of the form $(\mathbf{x}_k, -1)$; (ii) maximizes the minimum distance from the \mathbf{x}_i 's. This concept can then be extended to cope with non-linear separation rules and multi-classes classification tasks [32, Part II].

SVMs have already been exploited for building occupancy estimation tasks, e.g., in [17, 18]. The classical approach is to let \mathbf{x}_k contain functions of the current and past CO₂, temperature and ventilation levels (e.g., the average of $c(k), \dots, c(k-n)$). y_k is instead chosen to represent the building occupancy level $o(k)$. Once this mapping has been performed, it is possible to train a general multi-class SVM on the couples (\mathbf{x}_k, y_k) forming the training set. After this one can then estimate the unknown building occupancy by applying the estimated SVM on the \mathbf{x}_k forming the test set.

The SVM implemented in our tests that led to the best estimation error performance is a C-SVM exploiting a polynomial kernel of order 3. As features, it considers current and past values of the temperature, CO₂, and ventilation levels up to 1 hour in the past, and their first and second derivatives in time.

5.2 Neural Network (NN)

The Neural Networks (NNs) considered here correspond to maps of the form [33, Sec. 44]

$$y_k = \Psi'' \left(\sum_i \omega_i'' h_i(\mathbf{x}_k) + \theta'' \right), \quad h_i(\mathbf{x}_k) = \Psi'_i \left(\sum_j \omega_j' x_{k,j} + \theta_i' \right)$$

with y_k and \mathbf{x}_k having the same meanings of Section 5.1. The structure of the functions Ψ'' , Ψ'_i , h_i is a design parameter, but usually they remind how biological neurons electrically react to external stimuli. Once the design parameters have been chosen, training the network corresponds to search that particular set of weights \mathbf{w} for which the corresponding NN best fits the training examples. Once this function has been learned, it can be used for prediction purposes as did in the SVM case.

The NN implemented in our tests that led to the best estimation error performance is a complete feedforward network with Sigmoid activation rules and one hidden layer composed by 8 neurons. It considers the same features exploited to train the SVM based estimator.

6 Experiments

6.1 Description of the experimental setup

Tests have been performed in July 2013 in one of the rooms of the ACL-HVAC testbed, a fully instrumented facility located in the basement of the Q-building on KTH campus (see <http://hvac.ee.kth.se/> for more information). The information collected, available at <http://hvac.ee.kth.se/datasets.html>, comprises two weeks of measurements of CO₂ and temperature levels from HDH sensors, and of venting, cooling, and heating actuation levels from the PLCs controlling the IAQ of the room. Building occupancy

levels were manually registered for the whole period, with a time accuracy of 1 minute. To uniform the sampling times of the various signals (5 minutes), or in case of missing measurements, the information was resampled using linear interpolation schemes. The first week was used as a training set, while the second week was used as a test set.




Figure 1. The room where the dataset was collected. It is located in the basement of the Q-building on KTH campus, Stockholm, Sweden.




Figure 2. Schematic representation of the HVAC system for the control of the IAQ of the room shown in Figure 1.

6.2 Definition of the performance indexes

We consider four performance indexes:

- the *Mean Squared Error (MSE)* (12), characterizing the absolute estimation errors;
- the *accuracy* (14), reporting how many times the estimator returns the correct value;
- the *false positive / false negative occupancy detection rates* (17), describing the ability of discriminating the presence / absence of occupants in terms of false positives (when the room is estimated to be occupied while it is not) and false negatives (when the room is estimated to be empty while it is not).

More formally, let \mathbf{o} and $\hat{\mathbf{o}}$ be true and estimated realizations, respectively. Then the MSE associated to the couple $\mathbf{o}, \hat{\mathbf{o}}$ is simply

$$\text{MSE}(\hat{\mathbf{o}}) := \frac{1}{N} \sum_{k=1}^N (\hat{o}(k) - o(k))^2, \quad (12)$$

where $N = T_s$ in the offline estimator. To define the other performance indexes we then transform the signals \mathbf{o} , $\hat{\mathbf{o}}$ with codomain \mathbb{N}_+ (number of occupants) to signals with codomain $\{0, 1\}$ (room is non occupied, room is occupied) through indicator functions, i.e., through

$$\mathbb{1}(o(k)) := \begin{cases} 1 & \text{if } o(k) > 0 \\ 0 & \text{otherwise} \end{cases}, \quad \mathbb{1}(\mathbf{o}) := \begin{bmatrix} \mathbb{1}(o(1)) \\ \vdots \\ \mathbb{1}(o(N)) \end{bmatrix}. \quad (13)$$

Given (13), the accuracy of the estimate $\hat{\mathbf{o}}$ is

$$\text{Acc}(\hat{\mathbf{o}}) := \frac{N - \sum_{k=1}^N \mathbb{1}(o(k) - \hat{o}(k))}{N}. \quad (14)$$

To define the false positive / negative rates we introduce

$$\mathcal{N}_\theta := \{t \text{ s.t. } \mathbb{1}(o(k)) = \theta\}, \quad (15)$$

dividing the time indexes in two sets: \mathcal{N}_0 , for the k 's for which the room was not occupied, and \mathcal{N}_1 , for the k 's for which the room was occupied. With this it is possible to capture the mistakes “the room is estimated to be occupied while it is empty”, “the room is considered empty while it is occupied” with

$$\hat{\beta}(\theta) := \frac{1}{|\mathcal{N}_\theta|} \sum_{k \in \mathcal{N}_\theta} \mathbb{1}(\hat{o}(k)), \quad (16)$$

where we remark that the summation is performed over the set \mathcal{N}_θ . With (16) we then define the false positive and false negative rates as

$$\text{FP}(\hat{\mathbf{o}}) := \hat{\beta}(0), \quad \text{FN}(\hat{\mathbf{o}}) := 1 - \hat{\beta}(1). \quad (17)$$

6.3 Summary of the results

6.3.1 Comparison of the offline strategies

In this section we compare four estimators, all trained using the training set described in Section 6.1: the parametric and nonparametric strategies considered in Sections 3 and 4, indicated with the acronyms “P” and “NP”, and the machine-learning based algorithms of Section 5, indicated with the acronyms “SVM” and “NN”. Figure 3 shows the realizations of the estimates of the various strategies when applied to the corresponding test set, while Table 1 summarizes the achieved numerical performance.

We highlight the favorable properties of the deconvolution-based approaches with respect to the alternative methods considered above. We also notice that the performance of the parametric and nonparametric approaches are very close, fact that suggests that knowing accurately the physics of the room under consideration is eventually not required, as long as we consider building occupancy estimation problems. On the other hand, the parametric approach has generally smaller computational requirements, thus it seems appropriate to use prior information on the physics of the problem whenever this is available.

6.3.2 Evaluation of the sensitivity to the design parameter λ

As mentioned in Section 4.3, the regularization parameter λ enters in the design of the estimator by dictating the

| Algorithm | MSE | Accuracy | FP | FN |
|-----------|-------|----------|-------|-------|
| P | 0.114 | 0.880 | 0.008 | 0.018 |
| NP | 0.126 | 0.882 | 0.007 | 0.010 |
| SVM | 0.342 | 0.826 | 0.018 | 0.277 |
| NN | 0.268 | 0.811 | 0.067 | 0.095 |

Table 1. Summary of the performance achieved by the considered estimators in the test set.

typical variability of the estimated occupancy patterns. Our approach is to use, during the test phase, that particular λ^* that leads to the best estimation performance *in the training set*. Since the best λ in the test set may be different from the best λ in the training set, it is important to evaluate the effects of this unavoidable mismatch.

Figure 4 then plots the dependency of the MSE index on λ for both the parametric and nonparametric strategies and both the training and test set. It can then be noticed that the dependency on λ is weak: this is favorable, because it indicates that mismatches on λ are not going to disrupt the estimation performance.

6.3.3 Evaluation of the sensitivity to the design parameter N for the online estimation strategies

The length of the estimation window N plays also a design parameter role, since it trades-off computational requirements with information: the larger the window, indeed, the more holistic the deconvolution-based estimators are about the dynamics of the system. The intuition then suggests that there is a certain value for this horizon s.t. adding more information does not improve the estimation performance. This horizon is the one for that old dynamics do not influence the current estimates. The results shown in Figure 5 indicate that this length is, in our experiments, of about 5 days.

6.3.4 Evaluation of the importance of the knowledge of the actuation levels

When the occupancy levels are estimated using additional sensors that do not communicate with the central HVAC control systems, one has no information about how much air is injected in the rooms by the conditioning system. In our case, instead, we have access to this information: indeed our vision is that estimations should be performed by the control system itself, using its own information without exploiting additional sensors. A question is then: what is the value of the information on the HVAC system actuation levels? In our specific case (and without claims for generality), the answer is given in Table 2, reporting the performance indexes obtained by offline deconvolution-based estimators that neglect the signals $v(k)$. The outcome is that neglecting the ventilation levels lead to from two-fold to five-fold increases of the MSE, False Positive (FP) and False Negative (FN) levels, and a diminishing of approximately 6% of the accuracy of the estimates.

7 Conclusions

We proposed a scheme for the estimation of building occupancy levels in closed environments, exploiting information that is generally available to Heating, Ventilation and Air Conditioning (HVAC) controllers. The aim was to understand if such information is sufficient to perform mean-




Figure 3. Realizations of the estimation processes for the test set considered in our experiments.

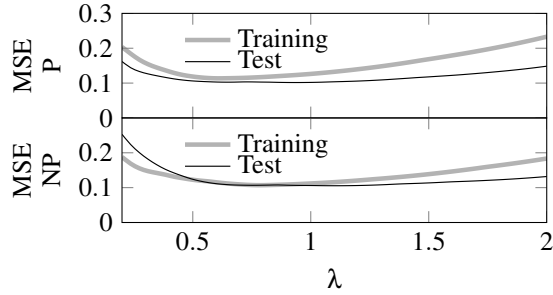


Figure 4. Sensitivity of the offline deconvolution-based estimators to the choice of the design parameter λ .

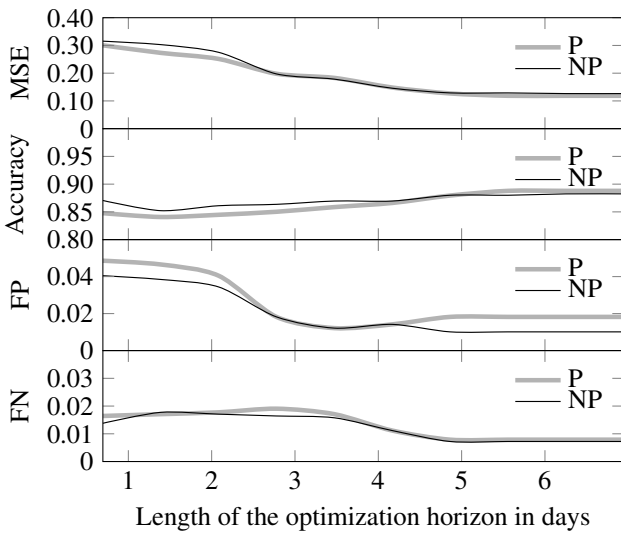


Figure 5. Dependency of the performance of the online estimators on the optimization interval length.

| Algorithm | MSE | Accuracy | FP | FN |
|-----------|-------|----------|-------|-------|
| P | 0.238 | 0.820 | 0.042 | 0.036 |
| NP | 0.235 | 0.821 | 0.034 | 0.045 |

Table 2. Summary of the estimation performance achieved by the deconvolution-based offline estimators that neglect the ventilation levels $v(k)$. Compare with Table 1 for a numerical evaluation of the importance of the knowledge of the air conditioning actuation levels.

ingful estimates of how the building occupancy changes in time.

The main assumption made is that the estimator can, for learning purposes and for a short period of time, access to direct measurements of the true occupancy levels. Another hypothesis is that the estimator uses the HVAC system actuation levels, in addition to the classical environmental signals such as CO₂ and temperatures.

The estimation strategy exploits dynamic Linear Time Invariant (LTI) models that are obtained by either parametric or nonparametric identification techniques. Based on the identified system, the room occupancy estimation problem is formulated as a regularized deconvolution problem that plugs prior information on the features of the building occupancy signal.

Numerically, the proposed algorithm returned the correct occupancy level in more than 88% of the times, reported the room as empty while it was occupied only in 1% of the times, and occupied while it was empty the 0.7% of the times. When compared with Support Vector Machine (SVM)-based and Neural Network (NN)-based estimators exploiting the very same information, our approach proved to have the best performance indexes.

The proposed technique was also used to evaluate the importance of considering HVAC system actuation levels for occupancy estimation purposes. We showed that, by neglecting these signals, we obtained a worsening of the estimates

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27

1 quantifiable in doubling the Mean Squared Error (MSE) in-
2 dexes and diminishing the accuracy by roughly 6%.

3 The idea considered in this paper can be extended towards
4 the construction of occupancy estimators for whole build-
5 ings, and thus for the identification of building occupancy
6 pattern models. Moreover, since the dynamics are assumed
7 linear, it may be possible to adapt the models identified in
8 a single room to other rooms of the same building, by an
9 opportune rescaling of the identified impulse responses ac-
10 counting variations in the structural properties of rooms.

11 Another appealing idea is to exploit blind system iden-
12 tification techniques to estimate both the system dynamics
13 and the building occupancy at the same time, thus removing
14 the assumption on the availability of the building occupancy
15 signal for a given period.

16 Acknowledgments

17 We thank Daniel Risberg, Giorgio Pattarello, and Lin Wei
18 for the technical support and help in collecting the dataset.

19 8 References

20 [1] F. Oldewurtel, D. Sturzenegger, and M. Morari, "Importance of
21 occupancy information for building climate control," *Applied Energy*,
22 vol. 101, pp. 521–532, Jan. 2013.

23 [2] D. Sturzenegger, F. Oldewurtel, and M. Morari, "Importance of Long-
24 Term Occupancy Information - A Validation with Real Occupancy
25 Data," in *Clima - RHEVA World Congress*, 2013.

26 [3] O. Guerra Santin, L. Itard, and H. Visscher, "The effect of occupancy
27 and building characteristics on energy use for space and water heating
28 in Dutch residential stock," *Energy and Buildings*, vol. 41, no. 11, pp.
29 1223–1232, Nov. 2009.

30 [4] Z. Yang, N. Li, B. Becerik-Gerber, and M. Orosz, "A multi-sensor
31 based occupancy estimation model for supporting demand driven
32 HVAC operations," in *Proceedings of the 2012 Symposium on*
33 *Simulation for Architecture and Urban Design*, 2012, pp. 98–105.

34 [5] S. Mamidi, Y.-H. Chang, and R. Maheswaran, "Improving building
35 energy efficiency with a network of sensing, learning and prediction
36 agents," in *Proceedings of the 11th International Conference*
37 *on Autonomous Agents and Multiagent Systems - Innovative*
38 *Applications*, 2012.

39 [6] N. Li, G. Calis, and B. Becerik-Gerber, "Measuring and monitoring
40 occupancy with an RFID based system for demand-driven HVAC
41 operations," *Automation in Construction*, vol. 24, pp. 89–99, July
42 2012.

43 [7] H. Wang, Q.-S. Jia, Y. Lei, Q. Zhao, and X. Guan, "Estimation of
44 occupant distribution by detecting the entrance and leaving events of
45 zones in building," in *2012 IEEE International Conference on*
46 *Multisensor Fusion and Integration for Intelligent Systems*. Ieee,
47 Sept. 2012, pp. 27–32.

48 [8] A. Schoofs, D. T. Delaney, G. M. P. O'Hare, and A. G. Ruzzelli,
49 "COPOLAN: non-invasive occupancy profiling for preliminary
50 assessment of HVAC fixed timing strategies," in *Proceedings of the*
51 *Third Workshop on Embedded Sensing Systems for Energy-Efficiency*
52 *in Buildings*, 2011.

53 [9] R. Melfi, B. Rosenblum, B. Nordman, and K. Christensen, "Mea-
54 suring building occupancy using existing network infrastructure," in
55 *2011 International Green Computing Conference and Workshops*,
56 July 2011, pp. 1–8.

57 [10] V. L. Erickson, Y. Lin, A. Kamthe, R. Brahme, A. Surana,
58 A. E. Cerpa, M. D. Sohn, and S. Narayanan, "Energy efficient
59 building environment control strategies using real-time occupancy
60 measurements," in *ACM Workshop on Embedded Sensing Systems for*
61 *Energy-Efficiency in Buildings*. New York, New York, USA: ACM
62 Press, Nov. 2009, p. 19.

63 [11] S. Meyn, A. Surana, Y. Lin, S. M. Oggianu, S. Narayanan, and T. A.
64 Frewen, "A Sensor-Utility-Network Method for Estimation of Occu-
65 pancy in Buildings," in *IEEE Conference on Decision and Control*,
66 2009, pp. 1494–1500.

[12] E. Hailemariam and R. Goldstein, "Real-time occupancy detection
67 using decision trees with multiple sensor types," in *Symposium on*
68 *Simulation for Architecture and Urban Design*, 2011.
69

[13] T. Leephakpreeda, R. Thitipatanapong, T. Grittiyachot, and
70 V. Yungchareon, "Occupancy-Based Control of Indoor Air Ventila-
71 tion: A Theoretical and Experimental Study," *Science Asia*, vol. 27,
72 no. 4, pp. 279–284, 2001.
73

[14] F. Scotton, L. Huang, S. A. Ahmadi, and B. Wahlberg, "Physics-based
74 Modeling and Identification for HVAC Systems," in *European Control*
75 *Conference*, 2013.
76

[15] S. Wang, J. Burnett, and H. Chong, "Experimental Validation of CO2-
77 Based Occupancy Detection for Demand-Controlled Ventilation,"
78 *Indoor and Built Environment*, vol. 8, no. 6, pp. 377–391, Nov. 1999.
79

[16] S. Kar and P. K. Varshney, "Accurate estimation of indoor occupancy
80 using gas sensors," in *International Conference on Intelligent Sensors,*
81 *Sensor Networks and Information Processing*. Ieee, Dec. 2009, pp.
82 355–360.
83

[17] K. P. Lam, M. Höyneck, B. Dong, B. Andrews, Y.-s. Chiou,
84 D. Benitez, and J. Choi, "Occupancy detection through an extensive
85 environmental sensor network in an open-plan office building," in
86 *IBPSA Conference*, 2009, pp. 1452–1459.
87

[18] B. Dong, B. Andrews, K. P. Lam, M. Höyneck, R. Zhang, Y.-S. Chiou,
88 and D. Benitez, "An information technology enabled sustainability
89 test-bed (ITEST) for occupancy detection through an environmental
90 sensing network," *Energy and Buildings*, vol. 42, no. 7, pp.
91 1038–1046, July 2010.
92

[19] G. Mustafaraj, J. Chen, and G. Lowry, "Development of room temper-
93 ature and relative humidity linear parametric models for an open office
94 using BMS data," *Energy and Buildings*, vol. 42, no. 3, pp. 348–356,
95 2010.
96

[20] D. Loveday and C. Craggs, "Stochastic modelling of temperatures
97 for a full-scale occupied building zone subject to natural random
98 influences," *Applied Energy*, vol. 45, no. 4, pp. 295–312, 1993.
99

[21] G. Lowry and M.-W. Lee, "Modelling the passive thermal response of
100 a building using sparse BMS data," *Applied Energy*, vol. 78, no. 1, pp.
101 53–62, 2004.
102

[22] J. C.-M. Yiu and S. Wang, "Multiple ARMAX modeling scheme for
103 forecasting air conditioning system performance," *Energy Conversion*
104 *and Management*, vol. 48, pp. 2276–2285, 2007.
105

[23] S. Wu and J.-Q. Sun, "A physics-based linear parametric model of
106 room temperature in office buildings," *Building and Environment*,
107 vol. 50, pp. 1–9, Apr. 2012.
108

[24] L. Ljung, *System identification - Theory for the user*, 2nd ed.
109 Prentice-Hall, 1999.
110

[25] G. Pillonetto, A. Chiuso, and G. De Nicolao, "Prediction error iden-
111 tification of linear systems: a nonparametric Gaussian regression ap-
112 proach," *Automatica*, vol. 47, no. 2, pp. 291–305, 2011.
113

[26] G. Pillonetto and G. De Nicolao, "A new kernel-based approach for
114 linear system identification," *Automatica*, vol. 46, no. 1, pp. 81–93,
115 2010.
116

[27] B. D. Anderson and J. B. Moore, *Optimal Filtering*. Prentice-Hall,
117 1979.
118

[28] G. Pillonetto and B. M. Bell, "Bayes and empirical Bayes semi-
119 blind deconvolution using eigenfunctions of a prior covariance,"
120 *Automatica*, vol. 43, no. 10, pp. 1698–1712, Oct. 2007.
121

[29] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical*
122 *Learning*. Springer New York Inc., 2001.
123

[30] R. Tibshirani and M. Saunders, "Sparsity and smoothness via the
124 fused lasso," *Journal of Royal Statistical Society: Series B (Statistical*
125 *Methodology)*, vol. 67, no. 1, pp. 91–108, 2005.
126

[31] A. Parisio, M. Molinari, D. Varagnolo, and K. H. Johansson, "A
127 Scenario-based Predictive Control Approach to Building HVAC Man-
128 agement Systems," in *IEEE International Conference on Automation*
129 *Science and Engineering*, 2013.
130

[32] B. Schölkopf and A. J. Smola, *Learning with kernels: Support vector*
131 *machines, regularization, optimization, and beyond*. The MIT Press,
132 2002.
133

[33] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation*
134 *Theory*. Prentice Hall, 1997.
135