Towards Smart Individual-Room Heating for Residential Buildings

Thomas von Bomhard[†], Dominic Wörner[*] and Marc Röschlin[*]
[*]Department of Management, Technology and Economics, ETH Zurich, Switzerland
dwoerner@ethz.ch, romarc@student.ethz.ch
[†]Institute of Technology Management, University of St. Gallen, Switzerland
thomas.vonbomhard@unisg.ch

***Abstract*—In many homes, residents keep their heating system always turned on although they are out or only occupy certain rooms, and thereby large amounts of energy are wasted. With our work, we aim to build an individual-room heating system that automatically detects occupancy, predicts a schedule based on that, and controls the heaters accordingly. First, we present our technical prototype for individual-room heating control. Second, we show that binary occupancy can be estimated using room climate sensors. We collected room climate data and occupancy data for three rooms over several days. We identified the relevant features and applied a Hidden Markov Model in a supervised and unsupervised way. We achieve a F1-score up to 85% for both variants in rooms which are occupied for longer periods. Third, we describe how a well-known occupancy prediction approach should be integrated into our heating control for optimal performance.**

***Keywords*—individual room heating, occupancy detection, hidden markov model, heating control**

I. Introduction

In Germany, more than 70% of energy consumed in residential buildings is due to space heating [1]. This distinguishes space heating as a major lever for energy-efficiency measures. Besides weather, thermal properties of the building, and the heating system itself, it is user behavior that determines the energy consumption of residential space heating. In this context user behavior can be summarized as follows: (1) setpoint preferences, (2) ventilation behavior and (3) usage of temperature setbacks in unoccupied rooms. Usually, adaptation of user behavior towards energy efficiency involves loss of convenience. For example, using temperature setbacks typically implies that residents have to remember to lower the temperature setpoint as they leave. Further, residents will face a cold room as they return.

In the 70s and 80s, programmable thermostats were developed to overcome the aforementioned issue. However, poor usability and the fact that they don't adapt to people's varying schedules led to misuse and their potential was never fully realized [2].

The ubiquity of smartphones triggered a new generation of Internet-connected thermostats which can be programmed remotely. This doesn't fix the problem that you have to remember to adapt the setpoints completely but it allows you to do it from wherever you are. Furthermore, big (touch-) screens and familiar user interfaces of smartphones make programming a schedule much easier.

Finally, intelligent thermostats promise to make programming schedules entirely obsolete. However, purely reactive systems which change schedules on arrival or departure analogous to automatic lighting are impractical because heating systems are slow in response. This is particularly true for hydronic heating systems which are standard in Germany. Therefore, the idea is to somehow predict the occupancy of a home or even individual rooms to control the heating accordingly. Already back in 1997, Mozer et al. [3] showed that even a highly nondeterministic schedule contains sufficient statistical regularity to be exploited by a predictive controller. But it took almost 15 years before scholars implemented protoypes and performed small scale field tests. Besides in-house occupancy detection based on motion sensors, door sensors etc. [4]–[6], also the geo-location feature of smartphones has been exploited for heating control [7]–[9]. In addition, implementations of both approaches are available as commercial products. The most prominent representative is NEST which is advertised as *the learning thermostat* and uses a passive infrared (PIR) motion detector to infer occupancy. The European counterpart tado relies on the geo-location feature and chooses a setpoint based on the residents current distance to their home and the time their home needs to heat up.

Thus far, little attention has been paid on individual-room solutions. Of the former, only [6] considered an individual-room HVAC system, but with reactive controls. In general, individual room-heating solutions allow for fine-grained controls which can be expected to provide further increase in energy-savings and convenience in contrast to a thermal thermostat. In order to provide individual-room heating controls every room has to be equipped with sensors and actuators. While a typical thermostat only carries a temperature sensor, we propose to measure the room climate along several dimensions. On the one hand this allows to give the residents detailed feedback about their room climate and allows to support efficient ventilation behavior. And on the other hand we will show that these room climate measurement data can contribute to determine the occupancy of individual rooms.

In this work, we present a prototypical system for individual-room heating control based on room climate sensors, controllable radiator valves and a smartphone app. The focus will be set to the most important *smart* feature: automatic occupancy-based heating. Our work is structured in the following way: First, we describe our prototypical system in Section II. Second, in Section III, we present a novel study which investigates occupancy detection using room climate sensor data by means of a Hidden Markov Model. Thereafter, the forthcoming implementation of an state-of-the-art occupancy prediction approach into our individual-room heating is explained in Section IV. We discuss limitations and advantages of our approach in Section V. We summarize our main findings and recommend for future work in Section VI.

II. System Overview

In the following, we briefly describe the components of our individual-room heating control system which is also illustrated in Figure 1a. For more details on the design and implementation from an information systems point of view, we refer to our work in [10].

A. Room Climate Sensors

We use off-the-shelf wireless Netatmo room climate sensors. The base station has the ability to measure temperature, relative humidity, barometric pressure, acoustics and CO_2 and is connected by wifi. Each base station can be extended by three additional room climate sensor units which lack the acoustics sensor. The communication between the base station and the additional sensor units is facilitated by a 868 Mhz radio. The sampling rate is fixed at 5 min and the measurement data is automatically sent to the Netatmo cloud where we can access it by an authenticated API. Consequently, we fetch the data and store it in our own database.

B. Backend

The Backend consists of a central server with a PostgreSQL database, a node.js application and a messaging broker. Furthermore, there are embedded Linux computers running an open source home automation software and another node.js application in every home. Those collect data from the controllable radiator valves and send control signals. The logic is implemented on the central server. Data and control signals are exchanged between the central server and the home controllers using a pub/sub architecture. Communication between the central server and the iPhone app is provided by an RESTful HTTP API.

C. Frontend

The main user interface of the system is an iPhone app which is shown in in Figure 1b. The user is able to get information about the current room climate as well as its history. Further, the user is able to control the heating by choosing temperature setpoints and by defining a schedule for each room individually. In addition the app makes use of the iPhone's geo-location capabilities in order to keep track of the user's distance to its home. There are additional features like status messages, displaying the expected heat-up time, feedback on air quality and optimal ventilation times considering current indoor and outdoor measurements.

D. Controllable Radiator Valves (CRV)

Controllable radiator valves are motorized valve heads with a radio module that allows for remote control. They can be installed in minutes by replacing the standard thermostatic valve heads at the wall-mounted radiators typically found in German homes.

III. Occupancy Detection Using Room Climate Sensors

A. Previous Work

Humans exhale moisture and CO_2 and also generate heat to some extent, which raises the interesting research question how good occupancy can be detected using room climate sensors. There are several approaches and evaluations [11]–[13] on solving occupancy detection using room climate but it was only investigated in office scenarios. Such a setting differs distinctively from the residential setting. While those offices were equipped with ventilation systems, a typical dwelling in central and northern Europe is ventilated manually by opening windows. Therefore, we did a first evaluation on occupancy detection using room climate sensors in residential buildings, which is described in [14]. We showed there that occupancy detection in that setting can be solved with accuracies well above 75%. We used the same occupancy data as in this work, but we only applied the HMM method in a supervised way. However, in a real-world setting, it is very unlikely that there will be such training data. In this work, we show a first approach to solve occupancy detection using an unsupervised method and compare it to the supervised method.

B. Data Acquisition

In order to collect occupancy ground truth we installed a camera in one of the apartments where our prototype is running. The camera was set up in the hallway monitoring the doors to three rooms. The apartment was a shared flat with two residents and three rooms: Their private rooms and the shared bathroom. We collected 15 days (14 in the case of bedroom 2) of continuous footage from 01/30/2014 until 02/13/2014. To avoid manually watching the entire video we used the open source computer vision library OpenCV [15] to extract sequences with movements. We then watched the snippets and labeled the events manually.

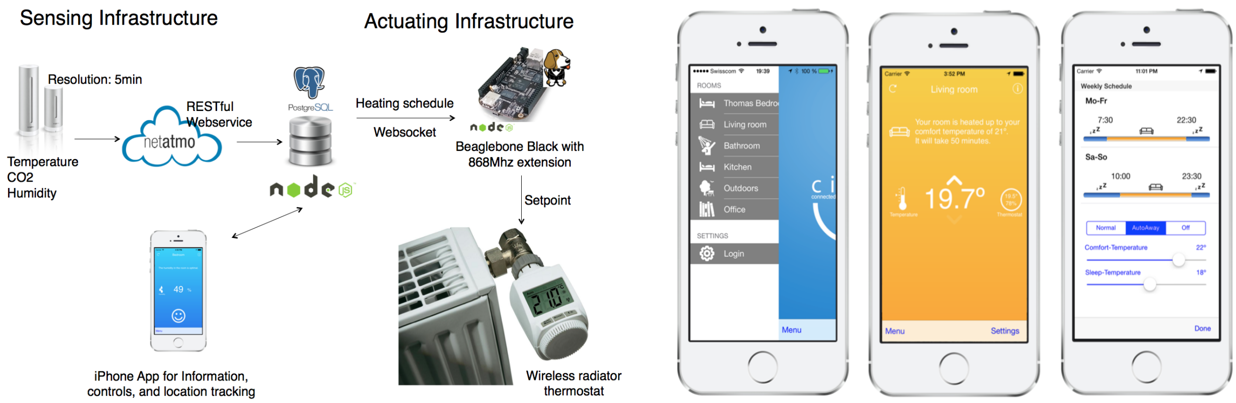

(a) System (b) App

Fig. 1. System overview and iPhone app

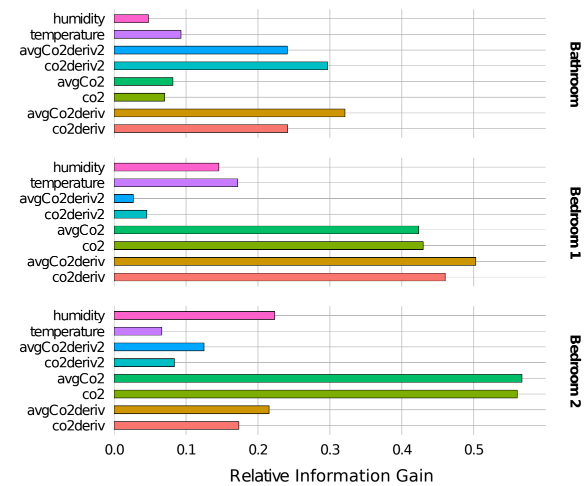

Fig. 2. Ranking in terms of relative information gain (RIG) for different features based on room climate sensor measurements. The entropy estimates are calculated by distributing the measurements to 30 equally spaced bins.

C. Methodology

1) Feature Selection: Given our room climate sensor measurement data we have to decide which features are useful to predict occupancy. The intuition is that the presence of people will affect room climate parameters. Living beings generate heat and exhale moisture and CO_2. We can therefore expect some coherence between these measures and occupancy. Intuitively, it seems legitimate to favor different features in each room, especially if rooms are used for different purposes and have different usage patterns. Therefore, we try to identify the features containing the most distinguishing power for each single room. In order to put the intuition on a more formal ground we consider the information theoretic notion of relative mutual information, also called relative information gain. It is defined by $RIG(x) = [H(y) - H(y|x)]/H(y)$ where H denotes the information theoretic entropy, x would be a feature, and y the random occupancy variable to be predicted. A more detailed discussion in a similar context can be found in [16]. To calculate the entropy of the room occupancy, as well as the entropy of the features under consideration we need the corresponding probability distributions $P(y)$ and $P(y|x)$. Since the feature distributions are not known apriori, we apply histogram-based entropy estimation. We use a total of 30 equally spaced bins while we discard two percent of the measurements at the boundaries of the distribution to increase robustness against outliers. These parameters have been evaluated by experimentation, although they are found to have marginal influence on the outcome of the entropy estimation.

The RIG of several features can be seen in Fig. 2. We note that features based on the CO_2 measurements lead to high RIG. However, there are obviously differences between the rooms. While the moving average of the first derivative of CO_2 scores highest in bedroom 1, it is rather low in bedroom 2. Moreover, the second derivative of CO_2 is irrelevant for both bedrooms, but relevant for the bathroom. For this reason, we rank the features according to their importance for each room individually. This will allow our predictive model to be tailored to each specific room and exploit the characteristics of the room climate to a greater extent.

The features we have chosen for each room are as follows:

Bathroom First derivative of CO_2, Second derivative of CO_2, Moving average of first derivative of CO_2

Bedroom 1 CO_2 concentration, First derivative of CO_2 concentration

Bedroom 2 CO_2 concentration, Moving average of CO_2 concentration

During the evaluation of feature candidates we did not only consulted the RIG in Fig. 2 but we also computed the RIG for groups of multiple features. Moreover, we took the pairwise correlation of different features into account. Fig.

2 only shows the informativeness of single features. Simply selecting the most informative features from Fig. 2 would not necessarily lead to the highest predictive power. Other feature combinations might complement each other in an even better way and reveal even more information about the occupancy of a room. Still, the RIG for single features is a very good guidance in feature evaluation.

2) Hidden Markov Model: A Hidden Markov Model (HMM) is a statistical model in which the dynamics are described by a discrete first-order Markov process with unobservable (hidden) states. The hidden states are assumed to generate a set of observable features X at every time step t. A HMM is specified by the transition probabilities $P(Q_{t+1} = q|Q_t = q_t)$ between subsequent hidden states q_{t+1} and q_t, by the emission distribution $P(X = x|Q = q)$ characterizing the features X, and by the initial state probabilities $P(Q_0 = q_0)$. We identify the hidden states to be either *unoccupied* or *occupied*, e.g. $q \in \{S_0, S_1\}$. The sequence $q_0, q_1, ..., q_T$ represents the binary occupancy of a room from time 0 to time T. If $q_t = S_1$ the room was occupied at time t and if $q_t = S_0$ the room was unoccupied, respectively. The observable features included in X are those selected in Sec. III-C1. They are processed in the form of one multidimensional vector x_t for every time step t. Each vector allocates one element per feature. The state-dependent emission distributions governing X are modeled by multivariate normal distributions.

3) Model Training and Validation: In general, the occupancy detection problem can be defined as a supervised learning problem or as an unsupervised learning problem. This is inherently different as in the first case the learning algorithm is allowed to use labeled data during the training phase which means the underlying occupancy sequence of the training data is given. In the unsupervised case the actual occupancy has to be estimated. The algorithm merely knows that there are two different hidden states producing the observable features.

To guarantee a sound training and validation phase we divide our data set into a training and a validation set. For the supervised approach we estimate the model parameters in a maximum likelihood fashion based on the training set. Then we compute the most probable sequence of occupancy states for both, the training and the validation set by applying the Viterbi algorithm [17], [18]. The performance on the training set provides information about how well the model fits the data in general whereas the performance on the validation set measures the actual classification capabilities on unseen data.

For the unsupervised HMM approach we use an expectation maximization algorithm [19] to estimate both, the model parameters and the hidden state sequence from the training data at the same time. Since the unsupervised training method does not know which of the two states represents an *occupied* or *unoccupied* room, we have to learn the mapping with the help of a heuristic after the prediction phase and apply it to the model's outcome accordingly. Similar to the supervised case we also evaluate the performance of the unsupervised approach on the validation data set, whereby the algorithm tries to predict room occupancy with the parameters learnt during training.

D. Results

1) Performance Metrics for Binary Occupancy Estimation: In order to evaluate the performance of the binary occupancy estimation the following metrics are considered: accuracy ($\frac{TP+TN}{TP+FP+TN+FN}$), precision ($\frac{TP}{TP+FP}$), sensitivity ($\frac{TP}{TP+FN}$), specificity ($\frac{TN}{FP+TN}$) and F1 score ($\frac{2TP}{2TP+FP+FN}$). Hereby, TP is the number of true positives, TN the number of true negatives, FP the number of false positives and FN the number of false negatives. Positive (negative) refers to the occupied (vacant) state. In all cases a higher performance metric means that the model is able to make a better prediction.

The motivation for this diversity of performance metrics is twofold. First, one might like to weigh wrong predictions differently. False negatives might be more severe than false positives as the user finds a cold room when it actually should be heated up. By reporting several performance metrics this distinguishability can be provided. Secondly, if a room is occupied or unoccupied most of the time the proportion of the two occupancy states is uneven. Even a trivial approach which simply always predicts either state can achieve spuriously good accuracies in such cases. Studying different performance metrics is crucial here, too.

2) Performance evaluation: In Fig. 3 we present the performance metrics that our supervised and unsupervised approach are able to achieve. For both bedrooms the supervised and the unsupervised approach reveal a remarkable distinguishing power. On unseen data from the validation set they both are able to reach a F1-score of over 85 % for bedroom 1. In the case of bedroom 2 the performance is diminished slightly but still achieves a considerable level of more than 63 %. We believe that for bedroom 2 our model was not able to grasp the relationship between occupancy and observable features in its entirety. This is due to the fact that the resident of bedroom 2 was not present in his room very often during our measurement phase and because he usually opened the window during night. CO_2 concentration is affected by this extraneous influence and complicates occupancy prediction.

Surprisingly, the unsupervised approach outperforms the supervised counterpart on the training data. On the one hand, the unsupervised approach indeed learns the underlying occupancy pattern very accurately which results in a robust prediction. But, on the other hand, we think that another validation is necessary when we have more data at hand for averaging out the performance figures.

For the bathroom we note that occupancy prediction is not viable in this form. The erratic usage pattern as well as the very short visits do not allow a meaningful inference on

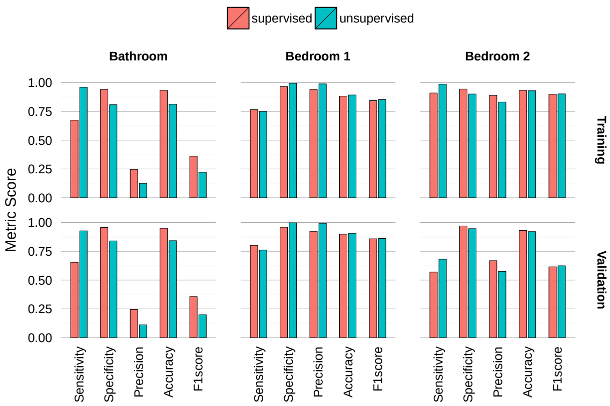

Fig. 3. Performance of our supervised and unsupervised HMM approach.

occupancy using our proposed model alone. On top of that, our sensors exhibit a time interval of five minutes between successive measurements. This makes it nearly impossible for the model to register such brief visits. Unfortunately though, they are the most prominent feature in the usage of a typical bathroom.

IV. Occupancy Prediction and Heating Control

Besides detecting whether or not a room is currently occupied, the system needs to be able to predict when a room will be (un)-occupied for the next hours and set the CRVs of the room accordingly.

A. Previous Work

There is vast amount of research on occupancy prediction algorithms for smart heating [4], [7]–[9], [20]. In 2013, *Kleiminger et. al* [21] presented a comparative performance analysis of state-of-the-art approaches based on a common dataset. It turns out that the "Presence Probabilities" approach by *Krumm et al.* [20] shows the best performance with 85% accuracy. Besides a few optimizations, the basic idea of that approach is to compute probabilistic schedule in which each 30-minute time slot of a day is given by the probability for the household to be *unoccupied* which is the ratio of the number of unoccupied occupancy states by the total number of occupancy states.

In addition to that, they combine their approach with an optimized version of the well-known "GPS-Controlled Thermostat" approach of Gupta et al. [7] which predict the arrival of residents based on the distance (or travel-to-home time) of them to their home. They use that information to reduce the thermostat setback temperature accordingly. *Krumm et al.* integrate that idea to their probabilistic schedule by always predicting *unoccupied* (even if the schedule suggests *occupied*) if a resident is outside a certain drive-time zone. They can show that this combined approach leads to a slight performance improvement which makes sense since it benefits from the advantages of the both algorithms. Finally, since this approach shows the best performance compared to other approaches, we will use and adapt it to our individual-room heating system.

B. Forthcoming Implementation

For each room, a probabilistic occupancy schedule is estimated using the room's past occupancy states, which were generated by our occupancy detection approach (as described in Section III). The heating controller checks constantly the schedule and adjusts the CRV for a given point in time as follows:
If it is *occupied*, the CRV holds the comfort temperature.
If it is *unoccupied*, the controller computes (1) the length of time to the next occupied interval I, denoted by t_{next}, (2) the necessary time to heat-up to comfort temperature, denoted by t_{heatup}, and (3) the travel-to-home time of the resident, denoted by $t_{traveltohome}$. Then the controller decides as follows: if $t_{next} > t_{heatup}$, the CRV is put to the setback temperature for saving energy. Otherwise, it checks if $t_{traveltohome} > t_{heatup}$. In that case, the CRV is put to the setback temperature and interval I of the schedule is set to *unoccupied* for that day. In the other case, the CRV is set to comfort temperature to preheat the home.

In summary, we apply Krumm et al. combined approach except that we only overrule the probabilistic schedule if a resident's travel time is higher than the heat-up time.

Our system estimates the heat-up time based on the average time of past heat-up procedures. Although this is a simple approximation, it reflects to some extent the specific building/room physics, the heating/radiator characteristics and the latest outdoor temperatures. Since our iPhone app sends significant location changes of a resident to the server, an estimate of the current travel-to-home-time can be easily calculated.

V. Discussion

A. Occupancy Detection

The results of our occupancy detection, as illustrated in Figure 3, indicate that an unsupervised HMM approach can be as good as the supervised HMM approach. For our aims of designing (and evaluating) a smart individual-room heating solution in practice, this is a great outcome. However, as already discussed in III-D, this needs further investigation using more data to improve reliability. Furthermore, with more data at hand, it would be interesting to evaluate the performance of applying a supervised HMM model, that was trained on a specific room type before, on a new room of similar type. This would help to overcome the training problem of the supervised approach in real-world settings.

B. Occupancy Prediction and Heating Control

By choosing one of the best occupancy prediction approaches, we expect also an high accuracy in our forthcoming evaluation. Moreover, since we will use the combined approach as described in Section IV-B, we expect higher energy savings without loss of comfort compared to systems that only use one method. The "Presence Probabilities" method helps to save more energy when residents are out-of-home but very close. In comparison to an approach that only regulates the heating based on the travel-to-home-time, because it will not decrease the temperature setpoint. On the contrary, this approach saves more energy in cases when residents do not come home as usual (and are not close to their home). In that case, an approach based only on "Presence Probabilities" heats up to comfort temperature according to the normal schedule. Consequently, the combined approach should have a better performance. However, in the case when a resident comes home irregular to the normal pattern it will still fail and there is no good solution for that. Nevertheless, the residents could take care of that, if they are provided with an intuitive control of the predicted schedule on their smartphone app.

VI. Conclusion and Future Work

In this work, we showed how an occupancy-based individual-room heating system using room climate sensors can be implemented for practical evaluation. More precisely, we presented: (1) the prototypical system for individual-room heating control, (2) an evaluation of how good one can infer occupancy from room climate data by using Hidden Markov Models in a supervised and unsupervised way, (3) a detailed explanation why and how a well-known occupancy prediction approach should be integrated into the heating control for optimal performance.

Our evaluation on occupancy detection has shown that a good performance (F1-Scores of about 85 % and 63 %) is achievable for rooms with longer occupancy periods. That is even the case, if an unsupervised HMM is applied. For rooms like the bathroom with short occupancy periods our solution can not be recommended.

Although the room climate sensor based approach has certain limitations, it is still beneficial for the purpose of occupancy-based heating. Most of the issues might be solved by considering additional measurements like motion or acoustics. Conversely, smart heating approaches based on motion detection could benefit from additional climate sensor-based occupancy estimation. Furthermore, since residents are interested in their room climate, such sensors become more widespread in homes anyway. For these reasons, we think that such a smart individual-room heating has great potential for saving energy, but this has to be evaluated by field experiment in future work.

References

[1] BMWi, *Bundesministerium für Wirtschaft und Technologie: Energiedaten nationale und internationale Entwicklung (Gesamtausgabe), BMWi Referat III C 3*, 2012.

[2] T. Peffer, M. Pritoni, A. Meier, C. Aragon, and D. Perry, "How people use thermostats in homes: A review," *Building and Environment*, vol. 46, no. 12, pp. 2529–2541, 2011.

[3] M. C. Mozer, L. Vidmar, R. H. Dodier *et al.*, "The neurothermostat: Predictive optimal control of residential heating systems," *Advances in Neural Information Processing Systems*, pp. 953–959, 1997.

[4] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse, "The smart thermostat: using occupancy sensors to save energy in homes," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems.* ACM, 2010, pp. 211–224.

[5] C. Ellis, J. Scott, M. Hazas, and J. Krumm, "Earlyoff: Using house cooling rates to save energy," in *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings.* ACM, 2012, pp. 39–41.

[6] T. Sookoor and K. Whitehouse, "RoomZoner : Occupancy-based Room-Level Zoning of a Centralized HVAC System," *Conference on Cyber-Physical Systems*, 2013.

[7] M. Gupta, S. S. Intille, and K. Larson, "Adding GPS-Control to Traditional Thermostats : An Exploration of Potential Energy Savings and Design Challenges," pp. 95–114, 2009.

[8] S. Lee, Y. Chon, Y. Kim, R. Ha, and H. Cha, "Occupancy prediction algorithms for thermostat control systems using mobile devices," *Smart Grid, IEEE Transactions on*, vol. 4, no. 3, pp. 1332–1340, Sept 2013.

[9] C. Koehler, B. D. Ziebart, J. Mankoff, and A. K. Dey, "Therml: Occupancy prediction for thermostat control," in *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ser. UbiComp '13. New York, NY, USA: ACM, 2013, pp. 103–112.

[10] T. von Bomhard, D. Wörner, and F. Wortmann, "Designing an information system for residential heating and ventilation to improve comfort and save energy," in *Advancing the Impact of Design Science: Moving from Theory to Practice*, ser. LNCS. Springer International Publishing, 2014, vol. 8463, pp. 398–402.

[11] K. P. Lam, M. Höynck, B. Dong, B. Andrews, Y.-S. Chiou, R. Zhang, D. Benitez, J. Choi *et al.*, "Occupancy detection through an extensive environmental sensor network in an open-plan office building," in *Proceedings of the 11th International IBPSA Conference*, 2009, pp. 1452–1459.

[12] Z. Han, R. Gao, and Z. Fan, "Occupancy and indoor environment quality sensing for smart buildings," in *Instrumentation and Measurement Technology Conference (I2MTC), 2012 IEEE International.* IEEE, 2012, pp. 882–887.

[13] E. Naghiyev, M. Gillott, and R. Wilson, "Three unobtrusive domestic occupancy measurement technologies under qualitative review," *Energy and Buildings*, vol. 69, no. 0, pp. 507 – 514, 2014.

[14] D. Wörner, T. von Bomhard, M. Röschlin, and F. Wortmann, "Look twice: Uncover hidden information in room climate sensor data," in *4th IEEE International Conference on the Internet of Things, IOT 2014, Cambridge, MA, USA, October 6-8, 2014.*

[15] G. Bradski, *Dr. Dobb's Journal of Software Tools.*

[16] B. Dong, B. Andrews, K. P. Lam, M. Hoeynck, R. Zhang, Y.-S. Chiou, and D. Benitez, "An information technology enabled sustainability test-bed (ITEST) for occupancy detection through an environmental sensing network," *ENERGY AND BUILDINGS*, vol. 42, no. 7, pp. 1038–1046, JUL 2010.

[17] A. J. Viterbi, "Error bounds for convolutional codes and an asymptotically optimum decoding algorithm," *Information Theory, IEEE Transactions on*, vol. 13, no. 2, pp. 260–269, 1967.

[18] G. D. Forney Jr, "The viterbi algorithm," *Proceedings of the IEEE*, vol. 61, no. 3, pp. 268–278, 1973.

[19] J. O'Connell and S. Højsgaard, "Hidden semi markov models for multiple observation sequences: The mhsmm package for R," *Journal of Statistical Software*, vol. 39, no. 4, pp. 1–22, 2011. [Online]. Available: http://www.jstatsoft.org/v39/i04/

[20] J. Krumm and A. B. Brush, "Learning time-based presence probabilities," in *Pervasive Computing.* Springer, 2011, pp. 79–96.

[21] W. Kleiminger, F. Mattern, and S. Santini, "Predicting household occupancy for smart heating control: A comparative performance analysis of state-of-the-art approaches. (submitted for publication)," 2013. [Online]. Available: http://www.vs.inf.ethz.ch/publ/papers/wilhelmk-predic-2013.pdf