

Building Relationships: Using Embedded Plug Load Sensors for Occupant Network Inference

Andrew J. Sonta^{1b} and Rishee K. Jain

Abstract—Understanding the underlying structure of building occupant dynamics is crucial to improving the effectiveness and energy efficiency of commercial buildings, as occupants fundamentally drive building design and operation. In current practice, we largely account for occupant behavior in the design and management of buildings through rudimentary schedules of presence or absence. However, the increasing availability of embedded sensors—such as plug load sensors—offers an opportunity not only to monitor occupants’ activity patterns but also to use these patterns to gain insight into the network structure of occupants. In this letter, we present a statistical methodology for inferring this network, which comprises social, spatial, and organizational ties among occupants. We apply our method to a 7-person office environment in Northern California, and we compare the inferred networks to ground truth social, spatial, and organizational networks obtained through validated survey questions. We demonstrate that this approach offers insights into the complex nature of occupant dynamics, which can ultimately serve as inputs into building design strategies that minimize energy consumption and improve occupant well-being.

Index Terms—Building occupants, buildings, inference algorithms, social networks.

I. INTRODUCTION AND RELATED WORK

BUILDINGS designed for commercial offices fundamentally exist to enable effective work, typically evaluated through measures of productivity, creativity, and/or collaboration. Recently, given the large environmental impact of buildings, building professionals have espoused energy efficiency as another key marker of a well-performing building.

The new paradigm of data analysis in the built environment—enabled by new urban embedded systems—has given researchers the opportunity to understand the operation of buildings and cities through these lenses of environmental performance and human activity (e.g., outdoor lighting control [1]; occupancy-driven operation of HVAC and lighting [2]). Few studies, however, have used sensor

data to model human activity patterns and the natural structure of occupant relationships in the built environment. As researchers in the field of organizational behavior have long noted, understanding these relationships can enable the design of new office layouts that improve workplace satisfaction and occupant performance [3]. Organizational relationships, or *ties*, are typically modeled through surveys or interviews that take considerable time and effort to administer. Often, ties are not measured at all, leaving managers with simple organizational charts that describe workforce breakdowns by department or project and lack any subtle insights into the true nature of workplace relationships. Understanding these relationships can also be useful for reducing energy consumption: Anderson *et al.* [4] found that social network structure largely affects the efficacy of eco-feedback campaigns, and we have found in previous work [5] that ties can be useful for coming up with new layouts that match occupant behavior with building systems.

While this previous work has noted the importance of understanding the occupant network for energy and occupant performance, little work has attempted to infer the true occupant network. Some statistical and data mining tools have been developed specifically for inferring network structure from time series data. These methods have typically been applied to biostatistical problems [6], though some recent work has considered the problem of inferring social networks from time series data about human activity [7]. In this letter, we expand upon our work in [8] and adopt network inference methods for the problem of inferring the occupant network structure from distributed plug load energy sensors—sensors which are becoming ubiquitous and, as discussed in [9], can be used to model individual activity states at the desk level.

Facility and organizational managers understand the importance of knowing the structure and performance of their organizations. While energy efficiency is becoming more of a priority, people and their productivity are still the most expensive pieces of an organization—for the University of California system, total employee salaries, wages, and benefits were roughly 74 times more expensive than utility bills. As a result, managers can be reluctant to make changes to their buildings if they worry about negative impacts to productivity. Research has long shown the potential energy efficiency impacts of making changes to building layouts, and recent studies have demonstrated that spatial changes in offices can improve employee well being and collaboration [3], [10]. In this letter, we aim to use sensor data to automatically infer the occupant network and understand space utilization

Manuscript received April 17, 2019; accepted August 18, 2019. Date of publication August 25, 2019; date of current version May 27, 2020. This work was supported in part by the Stanford Graduate Fellowship, in part by the Terman Faculty Fellowship, in part by the Center for Integrated Facility Engineering, and in part by the U.S. National Science Foundation under Grant 1461549 and Grant 1836995. This manuscript was recommended for publication by M. Gonzalez and P. Huang. (*Corresponding author: Andrew J. Sonta.*)

The authors are with the Department of Civil and Environmental Engineering, Stanford University, Stanford, CA 94305 USA (e-mail: asonta@stanford.edu; rishee.jain@stanford.edu).

Digital Object Identifier 10.1109/LES.2019.2937316

in more detail. Moving forward, we hope to enable work that can co-optimize for building and human systems that are fundamentally intertwined.

II. METHODOLOGY

In this section, we introduce a two-step process for inferring the occupant network from raw sensor data. The first step makes use of a method introduced in [9]: raw sensor data are collected from distributed plug load power strip sensors, and these data are mapped to abstracted states of occupant activities. In the second step, two different network inference models are adapted from the literature and introduced for estimating the network relationships among occupants as defined by their activities. We note that collecting and analyzing data related to occupant network relationships comes with several risks and ethical concerns if misused (e.g., loss of privacy and potential social embarrassment). In order to minimize these risks, we collected and analyzed data in accordance with the Institutional Research Board (IRB) and the IEEE Code of Ethics, which included creating a transparent process for obtaining consent for data collection, minimizing personal information collected, and automating anonymization of data during collection and analysis.

A. Determining Occupant Activities Through Plug Load Data

Consistent with our previous work [9], we define a time series of plug load energy use at the desk level for each occupant

$$\mathbf{X}_{i,d} = \{x_1, \dots, x_T\} \quad (1)$$

where i is the occupant index (for all occupants $1, \dots, I$), d is the day index (for all days $1, \dots, D$), and T is the total number of time steps in a single day (e.g., if data are collected at 15-min intervals, $T = 96$). For the full dataset, we complete a component selection process based on variational Bayesian inference to determine the number of activity states present in the plug load data. This enables our approach to be adaptable to differing number of natural states across study groups. An activity state can be defined as abstracted and categorized information describing occupant behavior based on plug load energy consumption. Once the number of activity states is inferred, we complete a classification process that ascribes each plug load energy datum to an activity state. The result is an abstracted time series that describes overall changes in activity states for all occupants in the study

$$\mathbf{X}_{i,d} \mapsto \mathbf{S}_{i,d} \quad (2)$$

where \mathbf{S} contains the activity states. At each time interval, all occupants are classified into one of the same number of states. For complete details on this state classification method, we refer the reader to [9]. In that work, we found that plug load sensors provide accuracy comparable to sensors specifically designed for occupant detection, and we have shown that shifts between states correspond accurately with actual changes in behavior such as going to a meeting.




Fig. 1. Influence model schematic with I entities (adapted from [7]), where S indicates state, and x indicates the signal.

B. Estimating the Occupant Network

Given time series data about occupant activities, the next step is to infer the occupant network as defined by relative similarities in the activity data. We define an occupant network as a graph $G = (\mathcal{V}, \mathbf{A})$, where \mathcal{V} is the set of occupants and \mathbf{A} is the adjacency matrix of the graph. We explore two options for inferring the adjacency matrix: the graphical lasso, which estimates the inverse covariance matrix [6]; and the influence model, which estimates an “influence matrix” [7].

Graphical Lasso: The graphical lasso was developed as a method for inferring sparse undirected graphical models—also known as Markov random fields—through L_1 (lasso) regularization. In [6], the data are defined as N multivariate observations with dimension p , mean μ , and covariance Σ . In our case, $N = D \cdot T$ (the total number of time steps), and $p = I$ (the number of workstations/individuals). The graphical lasso makes use of coordinate descent to estimate the inverse covariance matrix (Σ^{-1}), which is often considered as the adjacency matrix in a Markov random field.

Influence Model: The influence model, discussed in [7], models the interaction among entities quite differently. The model is based on a generally coupled hidden Markov model (HMM), in which the state of each entity (in our setting, S_i) at any given time point t is determined by the state of all entities $S_{1,\dots,I}$ at the previous time point $t-1$. A graphical representation of the influence model is shown in Fig. 1. The authors use expectation maximization to estimate the parameters of the model. One of the key parameters that is learned is the matrix \mathbf{R} —the influence matrix—often interpreted as an adjacency matrix in a network.

The output of each model is a matrix that we consider as the adjacency matrix defining a weighted, directed network. For consistency, we refer to these matrices as $\mathbf{A}^{\text{glasso}}$ and \mathbf{A}^{infl} , where the entry $A_{i,j}$ represents the strength of the tie from occupant i to occupant j . The two models proposed here for learning network structure are based on fundamentally different assumptions. In the graphical lasso, the entries of $\mathbf{A}^{\text{glasso}}$ can be interpreted as measures of conditional dependence (i.e., if the entry i, j is zero, entities i and j are conditionally independent given all other variables). In the influence model, the entries of \mathbf{A}^{infl} can be interpreted as the strength of influence between two entities given the time-step-dependent HMM assumptions embedded in the model. Fitting each of these models to the data may provide different insights (e.g.,

the graphical lasso may effectively model structural relationships, while the influence model may be better at capturing spontaneous trends in behavior, such as when two occupants take a break together).

III. RESULTS AND DISCUSSION

To analyze the performance of our occupant network inference methodology, we applied it to a dataset from a three-room, 7-person office setting in Northern California. We collected plug load energy consumption data at 15 m intervals through Onset HOBO data loggers for a two-week period at each of the 7 occupants' individual workstations. Each workstation included a computer and usually a monitor and other small office loads such as phone chargers. We applied the network inference method from Section II, producing two networks defined by adjacency matrices A^{glasso} and A^{infl} . We also collected ground truth spatial, social, and organizational network data through an online survey (discussed below) to benchmark the inferred network against validated methods for capturing strengths of relationships.

We assume that the ground truth network can be characterized by three equally important relationship components: spatial, social, and organizational. Previous work has suggested these three components are fundamental to the similarities and dissimilarities in occupant behavior [5], but future work should consider how each individual network component relates to the inferred network structure. We refer to each ground truth network component as A^{spatial} , A^{social} , and A^{org} .

Spatial Dimension: To embed the spatial dimension in the ground truth network, we set the entry $A_{ij}^{\text{spatial}} = A_{ji}^{\text{spatial}} = 1$ if occupants i and j are situated in the same office, and 0 otherwise.

Social Dimension: For the social dimension, we used the results from a survey question answered by all occupants. We adopted the "inclusion of the other in the self" scale [11], which has been shown to be effective in measuring the closeness of social relationships. In the survey, each occupant can choose any value between 1 and 7 to describe their perception of the closeness of their relationship to all other occupants. These values are then linearly scaled between 0 and 1, and they become the entries in the A^{social} matrix, where occupant i 's response about occupant j becomes the entry A_{ij}^{social} .

Organizational Dimension: We adopt the methods introduced in Krackhardt and Hanson [12] to measure the structure of the organizational network among the 7 occupants. The survey asks occupants whom else: 1) they share information with; 2) they seek technical advice from; and 3) they seek personal advice from. The relationship is interpreted to be 0 if none of the three are true and 1 if all three are true, with a linear scale between 0 and 1 if one or two of the characteristics of organizational relationships are true. Occupant i 's response about occupant j becomes the entry A_{ij}^{org} .

We combined these ground truth dimensions into a single network by equally weighting their adjacency matrices and adding them together. We refer to the combined network as A^* . Comparing the overall structure of two networks (such as A^* to A^{glasso}) remains a difficult problem in network science,

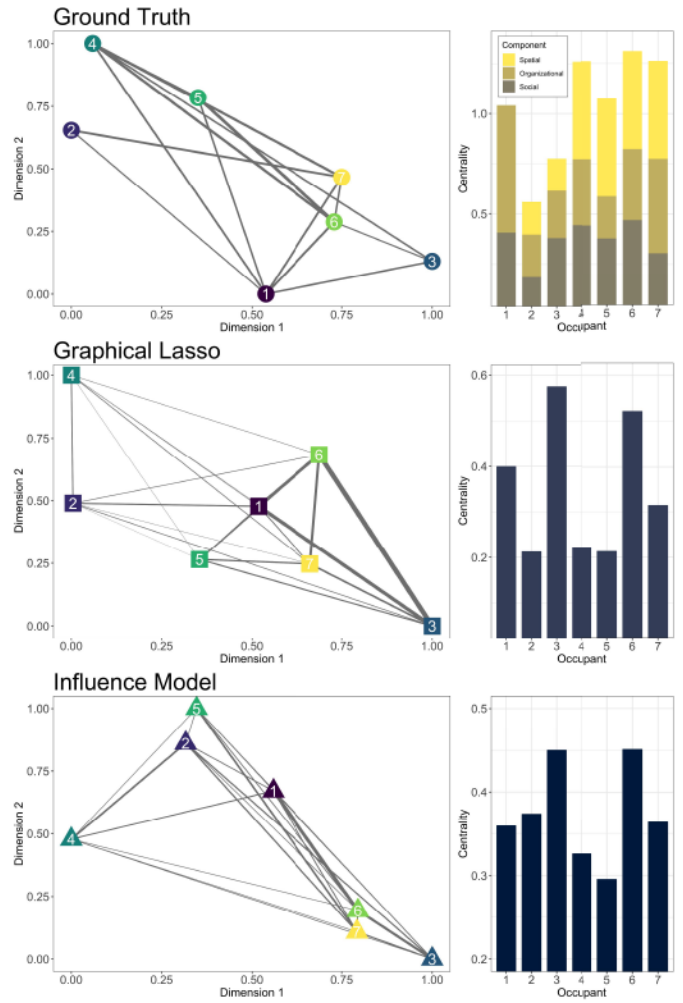


Fig. 2. Network comparison using *node2vec* representations in \mathbb{R}^2 (the two axes represent the 2-D). Degree centralities for each occupant shown to the right of each node embedding plot.

though certain common network-level techniques (e.g., community detection) can be used to compare high-level structure. To investigate the similarities and differences among our three networks, we adopt the *node2vec* algorithm introduced in [13]. The purpose of *node2vec* is to embed the network nodes into an n -dimensional feature space based on the topology and structure of the network, where n is chosen by the user. The algorithm can be useful when applying machine learning algorithms on network nodes, but it can also be used to understand underlying network structures and compare networks both visually and computationally. The smaller the Euclidean distance between two nodes in the embedded feature space, the more similar they are. The left-hand side of Fig. 2 shows the results of applying *node2vec* to the three networks (A^* , A^{glasso} , and A^{infl}). Here, we choose $n = 2$ for ease of plotting and visual comparison. Visually, the three networks all share a similar overall structure: occupants 1, 6, and 7 tend to be centrally located, with occupants 2, 4, and 5 on one side of them and occupant 3 on the other. The edges between the occupants indicate a relationship in the network, with the thickness of the line indicating relative tie strength. For someone with knowledge about this organization, it would make intuitive sense

that occupants 1 and 7 are centrally located: occupant 1 is the group director, and occupant 7 is the highest-ranking member. But the relative centrality of occupant 6 in the 2-D network representations requires further investigation.

To further explore these intricacies of network comparison, we plot the sum of in- and out-degree centralities for all three networks on the right-hand side of Fig. 2. In-degree centrality for node j in a network defined by adjacency matrix A is defined as $\sum_i A_{ij}$, and out-degree centrality for the same node is defined as $\sum_i A_{ji}$. We can see that the occupant 6 has large degree centrality resulting from all three components of the ground truth network, but especially the social component. This social centrality of occupant 6 highlights the value of the automated network inference method proposed in this letter. While the building or organizational manager might guess that occupant 1 (the director) or occupant 7 (the highest-ranking member) would be highly connected and relatively central occupants, he or she would have no reason to think that occupant 6 also has a central role without conducting a survey. When applied to large office environments (that could even involve multiple buildings), this type of network inference can offer subtle insights into social and organizational relationships that would require extensive time, effort, and investment to uncover through surveys.

Fig. 2 also provides subtle insights into the differences in the ways that the influence model and the graphical lasso analyze the data. Both occupants 1 and 6, according to the ground truth network, have large centrality from a social and organizational perspective. However, occupant 1 sits in his or her own office, and—based on the assumptions made about spatial relationships—has zero spatial centrality, whereas occupant 6 has high spatial centrality. It is interesting to note that occupant 6 has very high centrality in all three networks, whereas occupant 1 only has relatively high centrality in A^* and A^{glasso} . It is therefore possible that the influence model is more sensitive to spatial effects. As discussed above, the influence model may be more adept at capturing “spontaneous” changes in behavior, given its time-step-dependent HMM assumptions. It is quite possible that spatial cues are important for these kinds of spontaneous changes in behavior (e.g., one may be more likely to take a break if one sees one’s office-mate doing so). It is also interesting that occupant 3 has high degree centrality in both A^{glasso} and A^{infl} , but not in A^* —indicating that there may be other subtle aspects of relationships in activity patterns beyond the spatial, social, and organizational effects considered here.

IV. CONCLUSION

In this letter, we expand on a method originally introduced in [8] that automatically infers the *occupant network* from plug load energy consumption data collected at the desk level in an office building. By comparing inferred networks to the ground truth collected through surveys, we have shown that our method captures network characteristics that are both expected based on organizational structure and “hidden” in occupant social and spatial relationships. Through a case study, we discuss the similarities and differences of the two key

network construction methods—the graphical lasso and the influence model.

While our case study is small in size—and future work is required to demonstrate that valid inferences can be made in larger settings—it does demonstrate the potential for near real-time, automatic, and robust inference of the subtle aspects of organizational structure in buildings. In our future work, we aim to explore how spatial relationships obtained from floor plans can be embedded into the existing network inference methods, further adapting the inference to the specific context of building occupant networks. We also aim to explore how the networks change over time and in response to perturbations such as spatial rearrangement or the addition of a new employee.

In the end, a deeper understanding of building-occupant dynamics could enable the design and management of productive and energy-efficient spaces.

ACKNOWLEDGMENT

Any opinions, findings, and conclusions or recommendations expressed in this letter are those of the author(s) and do not necessarily reflect the views of National Science Foundation (NSF).

REFERENCES

- [1] Z. Kaleem, T. M. Yoon, and C. Lee, “Energy efficient outdoor light monitoring and control architecture using embedded system,” *IEEE Embedded Syst. Lett.*, vol. 8, no. 1, pp. 18–21, Mar. 2016.
- [2] A. Krioukov, S. Dawson-Haggerty, L. Lee, O. Rehmane, and D. Culler, “A living laboratory study in personalized automated lighting controls,” in *Proc. 3rd ACM Workshop Embedded Sens. Syst. Energy Efficiency Build. (BuildSys)*, 2011, pp. 1–6.
- [3] K. Sailer, R. Pomeroy, and R. Haslem, “Data-driven design—Using data on human behavior and spatial configuration to inform better workplace design,” *Corporate Real Estate J.*, vol. 4, no. 3, pp. 249–262, 2015.
- [4] K. Anderson, S. Lee, and C. Menassa, “Effect of social network type on building occupant energy use,” in *Proc. 4th ACM Workshop Embedded Sens. Syst. Energy Efficiency Build. (BuildSys)*, 2012, pp. 17–24.
- [5] A. J. Sonta, R. K. Jain, R. Gulbinas, J. M. F. Moura, and J. E. Taylor, “OESPG: Computational framework for multidimensional analysis of occupant energy use data in commercial buildings,” *J. Comput. Civil Eng.*, vol. 31, no. 4, Jul. 2017, Art. no. 04017017.
- [6] J. Friedman, T. Hastie, and R. Tibshirani, “Sparse inverse covariance estimation with the graphical lasso,” *Biostatistics*, vol. 9, no. 3, pp. 432–441, 2008.
- [7] W. Pan, W. Dong, M. Cebrian, T. Kim, J. H. Fowler, and A. S. S. Pentland, “Modeling dynamical influence in human interaction: Using data to make better inferences about influence within social systems,” *IEEE Signal Process. Mag.*, vol. 29, no. 2, pp. 77–86, Mar. 2012.
- [8] A. J. Sonta and R. K. Jain, “Inferring occupant ties: Automated inference of occupant network structure in commercial buildings,” in *Proc. 5th Conf. Syst. Built Environ.*, Shenzhen, China, 2018, pp. 126–129.
- [9] A. J. Sonta, P. E. Simmons, and R. K. Jain, “Understanding building occupant activities at scale: An integrated knowledge-based and data-driven approach,” *Adv. Eng. Informat.*, vol. 37, pp. 1–13, Aug. 2018.
- [10] F. Kabo, Y. Hwang, M. Levenstein, and J. Owen-Smith, “Shared paths to the lab: A sociospatial network analysis of collaboration,” *Environ. Behav.*, vol. 47, no. 1, pp. 57–84, Jan. 2015.
- [11] S. Gächter, C. Starmer, and F. Tufano, “Measuring the closeness of relationships: A comprehensive evaluation of the ‘inclusion of the other in the self’ scale,” *PLoS ONE*, vol. 10, no. 6, Jun. 2015, Art. no. e0129478.
- [12] D. Krackhardt and J. R. Hanson, “Informal networks: The company behind the chart,” *Harvard Bus. Rev.*, vol. 71, no. 4, pp. 104–111, Jul. 1993.
- [13] A. Grover and J. Leskovec, “node2vec: Scalable feature learning for networks,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Disc. Data Min.*, 2016, pp. 855–864.