OccuVAE: Integrating unsupervised occupancy inference in data-driven energy modeling for human-centric operation

Yufei ZHANG
yufei.zhang@epfl.ch
ETHOS Lab, School of Architecture, Civil and Environmental Engineering, EPF Lausanne
Switzerland

Andrew SONTA
andrew.sonta@epfl.ch
ETHOS Lab, School of Architecture, Civil and Environmental Engineering, EPF Lausanne
Switzerland

ABSTRACT

Data-driven building energy modeling is an emerging solution to facilitate the implementation of energy-flexible buildings. However, its black-box nature hinders interpretation, including with respect to human-building interaction. This drawback may bring risks to occupants' satisfaction under aggressive demand-side interventions. A modeling framework that successfully integrates occupancy inference with data-driven energy prediction can help to address these challenges without raising cost or privacy concerns. In this paper, we propose OccuVAE, which incorporates domain knowledge on human-building interaction into the black box of data-driven energy prediction, simultaneously inferring occupancy states from whole-building energy data. Its multifaceted capabilities are enabled by its architecture, consisting of both a Conditional Variational Autoencoder (CVAE), as well as an interpretable system sub-metering disaggregation module. We test OccuVAE on a synthetic office building subject to stochastic occupancy schedules and system operation. We demonstrate OccuVAE outperforms existing baselines for occupancy level extraction solely based on clustering energy-metering data (average F1 scores above 0.7 vs. baselines around 0.5). It also shows robust energy prediction performance for different prediction horizons while providing insights into system sub-metering disaggregation. We also demonstrate that it can recover occupancy level profiles from real-world energy use data of an office building, and we highlight necessary future steps to further address real-world challenges. This prototype is a critical first step toward holistic predictive operation leveraging both energy and occupancy flexibility.

CCS CONCEPTS

• **Applied computing** → *Engineering*.

KEYWORDS

Occupancy Level Detection, Building Energy Prediction, Demand-side Management (DSM), Variational Auto-encoder (VAE)

BuildSys '23, November 15–16, 2023, Istanbul, Turkey

ACM ISBN 979-8-4007-0230-3/23/11…$15.00
https://doi.org/10.1145/3600100.3626342

ACM Reference Format:
Yufei ZHANG and Andrew SONTA. 2023. OccuVAE: Integrating unsupervised occupancy inference in data-driven energy modeling for human-centric operation. In *The 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '23), November 15–16, 2023, Istanbul, Turkey*. ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3600100.3626342

1 INTRODUCTION

1.1 Background

The increasing penetration of intermittent renewables into the grid necessitates building energy flexibility and demand-side management (DSM), which requires building energy predictive models that easily scale to large and heterogeneous portfolios of buildings. In this context, data-driven building energy models are becoming promising alternatives to conventional physics-based models, due to abundant energy metering data in individual buildings and machine learning models' predictive performance [5]. However, the "black-box" nature of data-driven models often results in a lack of interpretability, including with respect to energy-related human-building interaction. Without explicit consideration of occupants, model-informed DSM interventions could have negative impacts on occupants' experiences such as comfort and satisfaction [9]. Occupants in office buildings are especially at risk, as the advances of automation in office buildings may inhibit occupant-centric energy management, such as when occupants are unable to override system operation when DSM measures prioritize energy use targets [15]. Therefore, it is an important step to ensure that occupancy information is known when analyzing and predicting building operation in the DSM context.

Information about occupancy in buildings is typically generated from targeted occupancy sensing systems, such as cameras and radio frequency identification (RFID) tags, which can help to consider the occupant experience in building management but may raise costs and privacy issues [16]. An alternative method leverages inverse modeling, which uses indoor environmental sensing signals, such as CO_2 concentration and humidity, to estimate occupant numbers [4, 7]. While this method offers high accuracy, its dependency on exhaustive data collection (including building thermal characteristics, system conditions, and occupant counts over a period of time) significantly hinders its scalability. At the same time, building-level energy data, which are more commonly available and fundamental to data-driven energy prediction in DSM, could also provide insights into occupant dynamics. One related research area is Non-Intrusive Load Monitoring (NILM), which utilizes data-driven models to disaggregate load metering data by appliances and can therefore offer insights into occupants' behavioral patterns. However, NILM methods typically rely on high-frequency load metering and require labeled datasets of appliance operation signatures [11]. As a result, they may not be suitable for engaging a large number of buildings that offer low-frequency flexibility through load shifting or shedding over minutes to hours. Prior studies also explored unsupervised extraction of occupancy levels and activity types from energy-metering data via latent variable models, (i.e., clustering-based models such as Gaussian Mixture Model (GMM) [18] and Hidden Markov Model (HMM) [6]). These methods work effectively in small residential buildings, since the magnitude and variance of the energy consumption is highly associated with the presence/absence of the occupants [2]. However, in larger and more complex buildings, such as office buildings, the interactions between humans and building systems drastically increase the difficulties in recovering occupancy status from building-level energy metering, due to the heterogeneity of occupants' behavior as well as the influence of building automation. Therefore, they still require detailed sub-metering (e.g., desk-level plug load) that is still unavailable in most of the buildings.

On the other hand, despite unavailable granular sub-metering, office-space energy consumption that is mainly driven by occupant behavior is metered separately from central plants in mechanical rooms (e.g., boilers and chillers) in most cases [12]. Prior studies have also identified pervasive operational patterns in each type of occupant-driven systems (e.g., lighting, miscellaneous electric loads (MEL)) [19] and quantitatively link the occupancy level to the energy consumption by simple functions [14]. This prior knowledge on human-building interactions can therefore be integrated into the data-driven energy modeling framework, especially the latent variable models that take into account underlying occupancy status. This will enable the disaggregation of system-level energy consumption from lumped building energy metering and ultimately produce underlying occupancy levels.

It is important to note that inferring occupancy level from whole-building office-space energy metering is a challenging task since such high-level energy metering cannot reflect all the variations in occupancy. For example, multiple occupants taking short trips away from the building may only lead to a slight drop in total energy consumption. Additionally, the quantitative relationships between occupancy levels and building system energy consumption in prior studies are empirical and simplified, which leads to errors when recovering both building system properties and underlying occupancy levels. Despite the potential for compromised accuracy in occupancy level inference, we believe this approach integrates valuable foundational insights into occupants' schedules into the data-driven building energy prediction workflow without raising costs and privacy issues due to additional occupancy sensing systems.

1.2 Main contributions

In this work, we propose OccuVAE, which enables unsupervised occupancy level estimation in the data-driven workflow of building energy prediction, given users' prior knowledge or assumptions on system operation modes. Its multi-faceted capability is enabled by two key design choices:

(1) A core Conditional Variational Autoencoder (CVAE) [17], a deep unsupervised latent variable model that can effectively generate high-fidelity complex data and simultaneously extract underlying influencing factors. Thus, it fundamentally enables inference of underlying occupancy levels during energy prediction.
(2) An interpretable load disaggregation network module inspired by recent advances in data-driven building energy modeling, which enforces consistency to prior knowledge on occupant-driven energy use in the model architecture [3], so the elements in the neural network explicitly associate to occupancy level, system capacities, etc.

Specifically, the proposed model provides original contributions by enabling:

- **An augmented data-driven building energy prediction pipeline with unsupervised occupancy level detection.** The model requires only commonly available whole-building energy metering in commercial buildings and is able to disaggregate system-level energy consumption from the lumped metering data to retrieve occupancy-level information.
- **An interpretable modeling framework for human-building interaction in commercial buildings through the introduction of prior knowledge.** This allows for customized definitions of the relationship between occupancy level and system energy consumption. The identified function parameters also reveal system properties (i.e., installed capacities).
- **An informative and interactive component for occupant-centric building management.** While we do not fully explore this capability in this initial work, it is a core motivation for the design of our model. In future work, it will enable probabilistic estimation of the occupancy level that helps quantify the risks of aggressive DSM measures affecting occupants' well-being. It will also allow for manual adjustment of the occupancy level profile when exploring the impacts of organizational schedule changes on energy consumption.

In this paper, we first explain the overall model architecture. Then, using a simulated case of a small office building subject to stochastic occupancy schedules and multiple occupant-centric operation strategies, while metered at the whole-building level. We benchmark its performance of energy prediction and occupancy level detection against other commonly used data-driven baselines. We also showcase the effectiveness of its occupancy level extraction capabilities from energy metering data of a real-world office building, as well as demonstrate necessary future steps to further address practical challenges.

2 MODEL OVERVIEW

OccuVAE follows an Encoder-Decoder structure with a latent space in the middle (Fig. 1). There are two pipelines in OccuVAE, representing two functionalities: **(1).** Predicting occupancy level and energy use for a future date (blue path in Fig. 1) based on historical energy metering data and easily-accessible future information such as site weather forecasting. **(2).** Inferring real-time occupancy level (red path in Fig. 1) from energy metering data at the current time. The inference pipeline is an autoencoder (AE) aiming to reconstruct identical input data during the training of the full model. Ideally, a model with minimized prediction and reconstruction error is capable of both generating target variables (i.e., load metering in our case) and extracting representative underlying factors (i.e., occupancy levels). A key design choice in OccuVAE was to design the entire model as a CVAE and introduce an interpretable system sub-metering disaggregation module into the decoder, which enabled us to model the occupancy level as a latent random variable and decouple it from other influencing factors of energy use profiles.

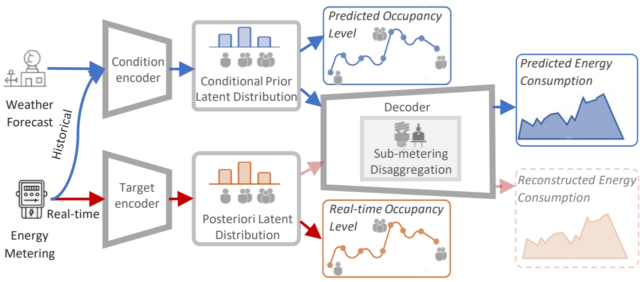

Figure 1: OccuVAE architecture. The prediction pipeline is in blue, while the real-time inference pipeline in red (dashed lines and light-color components only for training)

2.1 Design of CVAE

Here, we introduce the intuition behind the CVAE architecture. As an advanced deep unsupervised latent model, VAE endows a probability distribution into the latent space in a basic AE and learns a representative latent random variable z by reconstructing observed target Y. This is denoted as the inference process. CVAE [17] further considers a conditional generation process, which generates Y with only input condition X using an additional condition encoder. We chose CVAE as the model architecture because it is well-suited to our task from two main aspects: **(1)** The conditional generation process enables the prediction capability of our model (the prediction pipeline in Fig. 1), where Y is the observed total energy metering data, and X corresponds to predictive features such as historical energy metering and site weather forecasting. **(2)** The latent random variable z representing observed energy metering Y may reflect underlying occupancy levels, which need to be both predicted given X and inferred given actual Y. In our model, we use a categorical latent distribution representing several discrete occupancy levels like other clustering models. CVAE ensures that the training of prediction is guided by the original inference process that "sees" actual Y, guaranteeing stable training and fast convergence. This can be better explained with the loss function of CVAE (Eq. 1):

$$\mathcal{L}_{\text{CVAE}} = \alpha\underbrace{(\log p_\theta(\mathbf{Y} \mid \mathbf{z}) + D_{KL}(q_{\phi,infer}(\mathbf{z} \mid \mathbf{Y}) \| p_{\phi,pred}(\mathbf{z} \mid \mathbf{X}))}_{\text{Inference loss}}$$
$$+ (1-\alpha)\underbrace{(\log p_\theta(\mathbf{Y} \mid \mathbf{z}) + D_{KL}(p_{\phi,pred}(\mathbf{z} \mid \mathbf{X})) \| p(\mathbf{z}))}_{\text{Prediction loss}} \tag{1}$$

The total loss of CVAE is a weighted combination of prediction loss and inference loss. The log-likelihood terms stand for maximizing the likelihood of predicted and reconstructed Y. The Kullback–Leibler Divergence term (D_{KL}) in inference loss enforces the approximation of the predicted latent distribution $p_{\phi,pred}(\mathbf{z} \mid \mathbf{Y})$ (denoted as conditional prior) and the inferred latent distribution $q_{\phi,pred}(\mathbf{z} \mid \mathbf{X})$ (denoted as posterior). This enables $q_{\phi,pred}(\mathbf{z} \mid \mathbf{X})$ to implicitly supervise the predicted $p_{\phi,pred}(\mathbf{z} \mid \mathbf{Y})$. There is another D_{KL} term in prediction loss that enforces $p_{\phi,pred}(\mathbf{z} \mid \mathbf{Y})$ to approximate a given standard distribution $p(\mathbf{z})$ as the total prior of z (e.g., uniform distribution in the case of a categorical latent space). In office buildings, zero-occupancy hours usually dominate other occupied levels. We therefore only calculate D_{KL} for non-zero occupancy levels.

2.2 System sub-metering disaggregation

CVAE defines a principled latent space representing energy metering, which we hypothesize is related to occupancy levels. However, interactions between occupants and energy-intensive building systems are complex, which makes it difficult to draw a direct connection. To help make the link between the latent space and occupancy, we leverage prior work which has distilled these interactions into interpretable equations [14, 19]. In OccuVAE, we incorporate prior knowledge on energy-intensive human-building interaction through the following equation represented by one (bi-)linear layer in a neural network:

$$E_{sub-meter}(t) = P_{dynamic} \cdot z_{occ}(t) \cdot \Phi_{el,env}(c(t)) + P_{base} \tag{2}$$

Where $P_{dynamic}$ denotes the occupant-driven dynamic part and P_{base} an optional constant part in the total installed capacity for each type of system. As the main parameters to be fitted, their initialization can be based on reference values in standards or knowledge of the target building. This layer has the following inputs: **(1)**. $z_{occ}(t)$ refers to the time-varying occupancy level in a broad sense, which can be transformed from the latent occupancy level distribution (e.g., lighting roughly links to binarized occupancy, while MELs vary almost proportionally to a continuous occupancy rate within $[0, 1]$). By assigning a possible occupancy rate for each occupancy level, we can also calculate continuous expectation in $[0, 1]$ for the discrete distributions. This is a critical step to introduce prior knowledge of how occupants interact with different types of systems. By doing so, we are able to get more insights into the human-building systems (i.e., extracting continuous occupancy rates and even recovering sub-metering information from lumped metering). **(2)**. $\Phi_{el,env}(c(t))$ is an optional input (only for the bi-linear case), which denotes a possible discounting factor as a function of varying outdoor conditions $c(t)$ indicating potential passive strategies. This function can also be defined simply and fitted. For example, for daylighting, a linear layer with sigmoid activation can be fitted to solar irradiance, which indicates that only irradiance within a certain range reduces lighting energy. Finally, the decoder in Fig. 1 also contains a residual prediction module, which is a normal neural network capturing residual energy uses that are unable to be explicitly defined by the interpretable equations within the sub-metering disaggregation module.

3 PERFORMANCE BENCHMARKING

3.1 Synthetic office building operation data

To evaluate the performance of OccuVAE, we simulate synthetic operation data for a small office building based in San Francisco, which is subject to stochastic occupants' schedules as well as associated interactions with energy systems (following the workflow in [8]). We simulate the building's energy use for one year with hourly timesteps to obtain the synthetic building operation data. To match our problem statement, we aggregate energy use to the building level. Table 1 gives an overview of the synthetic building and the operation dataset.

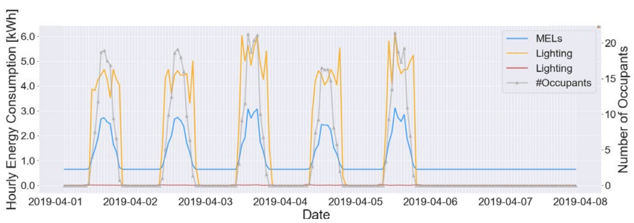

Figure 2: Exemplary weekly profiles of occupant counts and office-space energy use by systems in the synthetic building

Fig. 2 shows a typical weekly variation of occupants' counts and office-space energy use by systems. The profiles of occupant counts are stochastic, yet a slightly higher peak on Wednesday and Friday can usually be observed as meetings gather more people on these two days. As shown in the plot, MELs are proportionate with occupant counts given the direct link in the simulation setup. We note that it is common for MELs to be considered the best indicator of occupancy levels [19]. However, it is challenging to disaggregate this type of energy use from lumped energy metering dominated by lighting energy use. The lighting system is under DLR operation and therefore raises energy use in low-light hours, even though the occupancy level is usually quite low then. Finally, although we assume the terminal devices of HVAC (fan-coil units in this case) are also metered in the lumped energy metering, their magnitude is rather negligible in the total energy, as the most energy-consuming elements in HVAC systems such as chillers and air-handling units are metered separately.

3.2 Performance benchmarking setup

3.2.1 Configuration of OccuVAE. Following the procedure described above, we implement OccuVAE, which is composed only of multi-layer perceptrons (MLPs) without special consideration of temporal dependency. Our goal is to disaggregate MELs and lighting (with or without DLR) following the setups explained in Section 2.2, while leaving energy used by HVAC terminals for the residual prediction module, as their magnitudes are usually quite small. To approximate the continuous occupancy rate sweeping through $[0, 1]$, we found a four-class categorical distribution is sufficient (corresponding to four equally-spaced occupancy rates, namely $\{0, 1/3, 2/3, 1\}$). During training, for each hourly time step, the model does reconstruction and prediction at the same time, and occupancy levels are predicted and inferred without supervision, as explained in Section 2.1. We use the following input conditions: weather data of outdoor temperature and GHI, auxiliary time index variables (e.g., day of the week, hour of the day, etc.), and lagged targets (i.e., historical energy data lagged by hours or days, to enable processing of sequential data). Prediction horizons longer than one hour are also achievable by feeding back its previous prediction.

3.2.2 Baselines: Occupancy level extraction. We benchmark OccuVAE on both real-time inference and day-ahead prediction of occupancy levels. We select two clustering-based baselines: GMM and Input/Output HMM (IOHMM). GMM directly clusters hourly energy data with a mixture of Gaussian distributions, and it only handles real-time inference. IOHMM, a variant of HMM, models sequential dependencies of latent states and therefore handles both real-time inference and day-ahead prediction. It also accounts for external input variables (we consider weather conditions in this case) impacting the target energy metering. Further details of IOHMM can be found in [1]. For experiments, we used the open-source implementation of IOHMM in [13]. Unlike our model, the baselines do not infer continuous occupancy rates and are limited to three occupancy levels. We therefore use K-means clustering on the occupant counts to discretize actual continuous occupancy rates and that from OccuVAE to three levels in order to enable comparison, but note the strength of our model to infer continuous rates. We use the F1 Score to assess the models' performance on each occupancy level.

3.2.3 Baselines: Energy Prediction. We benchmark OccuVAE on both one-hour-ahead and day-ahead prediction, measured by Root Mean Square Error (RMSE). We select two widely-used baselines: Light Gradient Boosting Machine (LightGBM) and Long Short-Term Memory (LSTM). Both baselines are also trained for one-hour-ahead prediction with the same input conditions as that for OccuVAE. The only difference is that LSTM directly looks back to the sequential historical energy data instead of relying on lagged targets. To obtain day-ahead prediction, we also simply feedback on the one-hour-ahead prediction recursively for both baselines. In practice, there are better alternative setups for both LightGBM and LSTM to avoid error propagation in multi-step prediction, but here we only consider the day-ahead prediction as an empirical evaluation of robustness compared to our model, when similar training setup and features are provided to all the models.

It is also necessary to note that although our design for occupancy level extraction and system energy-use disaggregation is intended to produce more understanding of underlying human-building interaction, it may not guarantee better performance compared to "black-box" baselines of data-driven energy prediction. It is possible that the interpretable components may even constrain the solution space of model architecture optimization and end up with compromised prediction capability. Nevertheless, comparing the advantages and disadvantages of our model against other established baselines will help us understand how to address the compromised prediction performance in future works.

We also demonstrate OccuVAE's unique capability to recover dynamic sub-meter load capacity to further provide empirical insights into its advantages and drawbacks.

Table 1: Building condition summary and dataset overview for the synthetic building

Building condition and system operation	
Location	San Francisco, California, USA
Area	564 m^2
Occupants	Two occupant groups: - Regular group: 9 am-6 pm (variation: 20 mins) - Flexible group: 11 am-4 pm (variation: 60 mins) Events: meeting on Wed. and Fri., most likely 60 mins
Lighting	On/off control in each room, switch on when at least one occupant, together with daylight-responsive (DLR) diming in daylit area
Equipment (MELs)	Proportionally-changed energy use with occupant counts
HVAC	Normally-distributed stochastic room temperature setpoints
Dataset overview	
Accessible energy metering	Lumped metering occupant-driven end-use systems in office space, including lighting, MELs, HVAC terminal units (Fan-coil units in this case), separated from the metering of centralized systems (chillers, air handling units)
Site weather data	Outdoor air temperature and global horizontal irradiance (GHI) from the weather file
Temporal granularity	Hourly
Data Split Periods	• Training period: 01/01-01/07 • Validation period: 01/07-01/08 • Testing period: 01/08-31/12

Table 2: Summary of performance benchmarking

Occupancy Level Extraction (F1 Score)									Energy Prediction (RMSE)		
Models	Day-ahead Prediction				Real-time Inference				Models	Day-ahead	Hour-ahead
	Low	Med.	High	Avg.	Low	Med.	High	Avg.			
OccuVAE	0.96	0.44	0.74	0.71	0.97	0.48	0.76	0.73	OccuVAE	0.79	0.48
IOHMM	0.82	0	0.55	0.46	0.94	0	0.55	0.50	LightGBM	1.61	0.43
GMM	-				0.95	0.01	0.62	0.52	LSTM	0.80	0.49

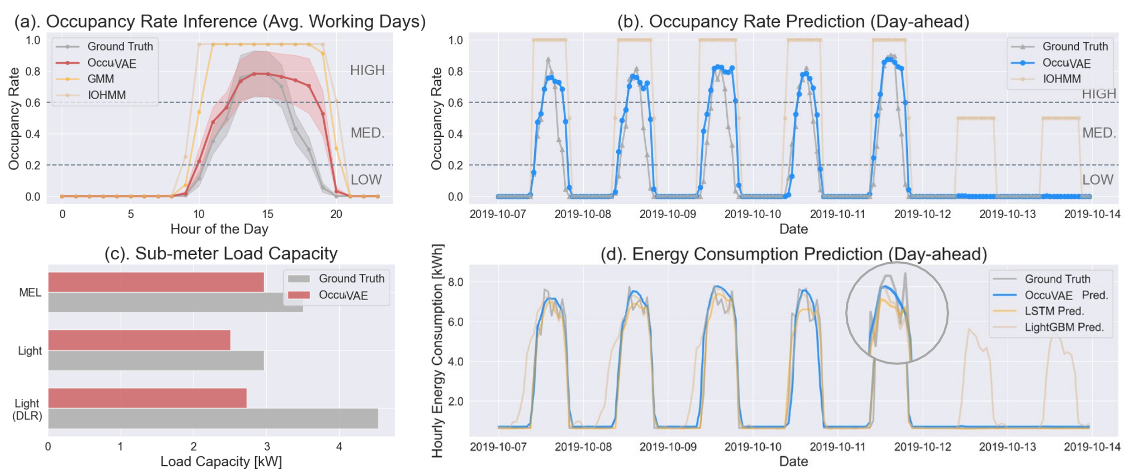

Figure 3: Performance demonstration: (a) Averaged daily occupancy rate real-time inference (with hourly one-standard-error). (b) Exemplary week of day-ahead occupancy rate prediction (ranges of discrete occupancy levels labeled). (c) Identified sub-meter load capacity. (d) Exemplary week of day-ahead energy prediction (with zoom-in for peak-demand hours)

3.3 Results of performance benchmarking

3.3.1 Occupancy Level Extraction. As shown in Table 2, for both day-ahead prediction and real-time inference, OccuVAE attains average F1 scores above 0.7, with both baselines near 0.5. Additionally, observing Fig. 3.(a), OccuVAE captures the distribution of actual peak occupancy rates very well. We believe this is partly enabled by the sub-metering disaggregation module that recovers MELs from the total energy metering (as shown in Fig. 3.(c)). Without the explicit sub-metering disaggregation framework, the baselines infer very high occupancy levels during all occupied hours, likely as a result of lighting energy. OccuVAE more accurately models the ramping-up and -down times, but we note that OccuVAE presents a mismatch with the ground truth during the afternoon medium occupancy hours, when energy levels are higher as occupancy has dropped (as Fig. 3.(d) shows). This is likely due to the DLR operation, with daylight being less abundant in the late afternoon, resulting in increased energy consumption. We note that in terms of sub-metering disaggregation, OccuVAE shows the largest mismatch with DLR-lighting (Fig. 3.(c)). This suggests the DLR mechanism is oversimplified in the current model and could be improved in future studies.

3.3.2 Energy Prediction. As shown in Table 2, all the models perform similarly on one-hour-ahead prediction. Their performances are compromised inevitably in day-ahead prediction due to potential error propagation caused by the recursive feedback prediction, especially for LightGBM. As Fig. 3.(d) shows, without special design for multi-step prediction, LightGBM struggles and fails to distinguish non-working days, while OccuVAE and LSTM remain robust. However, looking closely at peak demand hours from Fig. 3.(d), we see OccuVAE struggling to capture detailed patterns in energy use. Aside from the challenges with modeling DLR mentioned above, this mismatch may also reflect potential conflicts between the sub-metering disaggregation module that is constrained by interpretable functions linking occupancy level and energy use and the residual prediction module that only aims for accurate prediction. We expect to address this issue in future work by offering special attention to peak-demand hours and introducing additional regularization on module parameters.

4 REAL-WORLD CASE STUDY

We also present the initial results of unsupervised occupancy level prediction and detection on a real-world office space operation dataset. The target building is Building 59 of the Lawrence Berkeley National Laboratory (LBNL-Bldg 59) [10]. This dataset contains system-level energy metering, occupant counting, HVAC operation, as well as indoor and outdoor environmental conditions. Data was collected over three consecutive years (2018-2020) and witnessed several periods of unusual events, from evacuation during wildfire to lockdown due to COVID-19. For this initial work, we use an interval under stable operating conditions (05/2018 – 02/2019), and we empirically demonstrate our model's capability as well as shortcomings in occupancy level prediction and detection.

Table 3 gives an overview of the building and system conditions. All the information is drawn from the dataset curation report [10], which is also a typical source in practice to obtain prior knowledge of system operation patterns for subsequent modeling. The

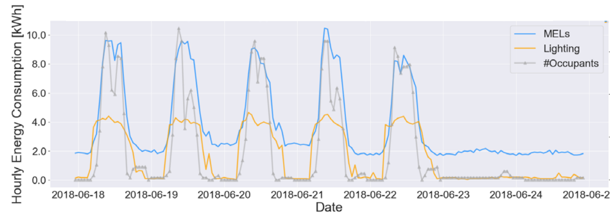

Figure 4: Examplary weekly profiles of occupant counts and office-space energy use by systems in the LBNL-Bldg 59

end-use systems in office space and the centralized HVAC systems are connected to two separate switchboards and metering systems in LBNL-Bldg 59. Although sub-metering by systems is also provided, here we follow the general assumption that only aggregated building-level energy metering is available. Fig. 4 shows a typical weekly variation of occupants' counts and office-space energy use by systems. Compared to the synthetic building in section 3.1, the lighting energy consumption is no longer dominating. It also shows a more stable pattern as it is controlled by occupancy sensors and is not subject to daylight conditions. The MELs are a larger proportion of overall load compared to the synthetic building. Additionally, unlike the synthetic case, MELs do not follow occupant counts exactly and vary more smoothly, particularly during midday lunch hours.

In this case study, we sub-sample the raw data to hourly granularity and use the identical configuration of OccuVAE as explained in section 3.2.1, except that we remove the lighting component driven by DLR from the load disaggregation module according to the building system condition. As shown in Fig. 5(a) and (b), the general patterns in the daily variation of the occupancy level as well as the peak hours and peak occupancy rates were identified in real-time inference. However, the moderate variation of MELs does lead to an overestimated occupancy rate profile, especially during lunchtime. It can also be observed that the inferred occupancy rate profile shows less variation, both hourly (Fig. 5a) and daily (Fig. 5b). This shortcoming also hinders the performance of occupancy rate prediction. As explained in section 2.1, the prediction pipeline is "supervised" by the inference pipeline. Although the predicted profiles are able to follow the inferred profiles, they are not consistent with the actual salient daily variation and end up with similar prediction results each day.

The case study on LBNL-Bldg 59 demonstrates the fundamental challenge that there can exist systematic error when extracting occupancy level from building-level energy metering, given the simplified relationships between energy and occupancy embedded in prior knowledge. Future improvements in the model may minimize this systematic error. An interesting direction of future work could be calibrating the extracted occupancy level through other data sources (e.g., surveys of working preferences). Still, the case study demonstrates that some level of human behavior information can be recovered in a privacy preserving manner at the whole-building level, and which could be used in future work focused on human-centric building operation in DSM scenarios.

Table 3: Building condition summary and dataset overview for LBNL-Bldg 59

Building condition and system operation	
Location	Berkeley, California, USA
Office-space area	Around 1600 m^2 (two office floors of the south-wing building)
Lighting	On/off control by occupancy sensors in each lighting zone
Equipment (MELs)	Plug-in by occupants
Dataset overview	
Accessible energy metering	Lumped metering occupant-driven end-use systems in office space, including lighting and MELs, separated from metering of centralized systems (heap pumps and root-top air units)
Site weather data	Outdoor air temperature and GHI from the site weather measurements
Temporal granularity	Sub-sampled to hourly
Data Split Periods	• Training period: 05/2018-11/2018 • Validation period: 11/2018-12/2018 • Testing period: 12/2018-02/2019

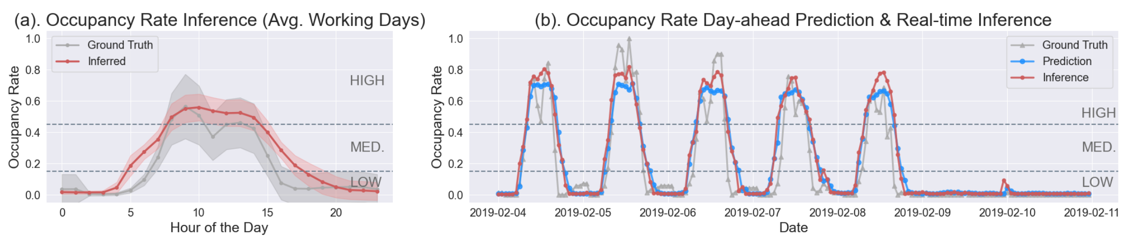

Figure 5: Performance demonstration on LBNL-Bldg 59: (a) Averaged daily occupancy rate real-time inference (with hourly one-standard-error). (b) Exemplary week of day-ahead occupancy rate prediction and real-time inference (ranges of discrete occupancy levels labeled, obtained from clustering on actual occupant counts).

5 DISCUSSION AND CONCULUSION

In this paper, we proposed OccuVAE, which leverages domain knowledge on human-system interaction alongside a novel CVAE-based system architecture. This approach enabled the integration of occupancy inference with data-driven energy prediction at the building level, performing well compared to the baselines in occupancy inference, occupancy prediction, and energy prediction while maintaining interpretability. OccuVAE also recovers system sub-metering information from lumped energy-metering, though further tests are required to benchmark this feature with real building data.

This integrated framework could be used as a tool in the context of occupant-centric system operation at the building level. For example, when the load aggregator or the building system operator plans DSM actions of HVAC operation, such as temperature setback or pre-heating, it can offer probabilistic prediction of the occupancy level that helps quantify the risks of affecting occupants' comfort. Most importantly, this prototype is a critical first step toward holistic and integrated occupant-centric management of energy systems alongside human and organizational systems. For example, this work could serve as a tool that allows manual adjustment of the occupancy level profile, therefore revealing additional energy flexibility opportunities when occupancy-flexible arrangements are considered (e.g., hybrid working or shared office).

While this work is preliminary, we expect to address shortcomings in future work—including better representation of interactions among human, building, and outdoor environments in the sub-metering disaggregation modules, as well as how to address its conflicts with the residual prediction modules. We plan to extend this modeling framework so the extracted occupancy level can be further calibrated through other data sources (e.g., surveys of working preferences). We also plan to further benchmark OccuVAE's performance on more real-world building data. In the end, OccuVAE is a step toward human-centric building analysis and operation as our buildings and grids undergo rapid change.

REFERENCES

[1] Y. Bengio and P. Frasconi. 1996. Input-Output HMMs for Sequence Processing. *IEEE Transactions on Neural Networks* 7, 5 (Sept. 1996), 1231–1249. https://doi.org/10.1109/72.536317
[2] Dong Chen, Sean Barker, Adarsh Subbaswamy, David Irwin, and Prashant Shenoy. 2013. Non-Intrusive Occupancy Monitoring Using Smart Meters. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys'13)*. Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/2528282.2528294
[3] Loris Di Natale, Bratislav Svetozarevic, Philipp Heer, and Colin N. Jones. 2022. Physically Consistent Neural Networks for Building Thermal Modeling: Theory and Analysis. https://doi.org/10.48550/arXiv.2112.03212 arXiv:2112.03212 [cs, eess]
[4] Afrooz Ebadat, Giulio Bottegal, Damiano Varagnolo, Bo Wahlberg, and Karl H. Johansson. 2013. Estimation of Building Occupancy Levels through Environmental Signals Deconvolution. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys'13)*. Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/2528282.2528290
[5] Anjukan Kathirgamanathan, Mattia De Rosa, Eleni Mangina, and Donal P. Finn. 2021. Data-Driven Predictive Control for Unlocking Building Energy Flexibility: A Review. *Renewable and Sustainable Energy Reviews* 135 (Jan. 2021), 110120. https://doi.org/10.1016/j.rser.2020.110120

[6] Wilhelm Kleiminger, Christian Beckel, Thorsten Staake, and Silvia Santini. 2013. Occupancy Detection from Electricity Consumption Data. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys'13)*. Association for Computing Machinery, New York, NY, USA, 1–8. https://doi.org/10.1145/2528282.2528295

[7] Han Li, Tianzhen Hong, and Marina Sofos. 2019. An Inverse Approach to Solving Zone Air Infiltration Rate and People Count Using Indoor Environmental Sensor Data. *Energy and Buildings* 198 (Sept. 2019), 228–242. https://doi.org/10.1016/j.enbuild.2019.06.008

[8] Han Li, Zhe Wang, and Tianzhen Hong. 2021. A Synthetic Building Operation Dataset. *Scientific Data* 8, 1 (Aug. 2021), 213. https://doi.org/10.1038/s41597-021-00989-6

[9] Jingjing Liu, Rongxin Yin, Lili Yu, Mary Ann Piette, Marco Pritoni, Armando Casillas, Jiarong Xie, Tianzhen Hong, Monica Neukomm, and Peter Schwartz. 2022. Defining and Applying an Electricity Demand Flexibility Benchmarking Metrics Framework for Grid-Interactive Efficient Commercial Buildings. *Advances in Applied Energy* 8 (Dec. 2022), 100107. https://doi.org/10.1016/j.adapen.2022.100107

[10] Na Luo, Zhe Wang, David Blum, Christopher Weyandt, Norman Bourassa, Mary Ann Piette, and Tianzhen Hong. 2022. A Three-Year Dataset Supporting Research on Building Energy Management and Occupancy Analytics. *Scientific Data* 9, 1 (April 2022), 156. https://doi.org/10.1038/s41597-022-01257-x

[11] Alan Meier and Dan Cautley. 2021. Practical Limits to the Use of Non-Intrusive Load Monitoring in Commercial Buildings. *Energy and Buildings* 251 (Nov. 2021), 111308. https://doi.org/10.1016/j.enbuild.2021.111308

[12] Clayton Miller, Anjukan Kathirgamanathan, Bianca Picchetti, Pandarasamy Arjunan, June Young Park, Zoltan Nagy, Paul Raftery, Brodie W. Hobson, Zixiao Shi, and Forrest Meggers. 2020. The Building Data Genome Project 2, Energy Meter Data from the ASHRAE Great Energy Predictor III Competition. *Scientific Data* 7, 1 (Oct. 2020), 368. https://doi.org/10.1038/s41597-020-00712-x

[13] Mogeng, Thuener Silva, and Eric Denovellis. 2023. Input Output Hidden Markov Model (IOHMM) in Python. https://github.com/Mogeng/IOHMM.

[14] Martín Mosteiro-Romero, Clayton Miller, Adrian Chong, and Rudi Stouffs. 2023. Elastic Buildings: Calibrated District-Scale Simulation of Occupant-Flexible Campus Operation for Hybrid Work Optimization. *Building and Environment* 237 (June 2023), 110318. https://doi.org/10.1016/j.buildenv.2023.110318

[15] Zoltan Nagy, Burak Gunay, Clayton Miller, Jakob Hahn, Mohamed M. Ouf, Seungjae Lee, Brodie W. Hobson, Tareq Abuimara, Karol Bandurski, Maíra André, Clara-Larissa Lorenz, Sarah Crosby, Bing Dong, Zixin Jiang, Yuzhen Peng, Matteo Favero, June Young Park, Kingsley Nweye, Pedram Nojedehi, Helen Stopps, Lucile Sarran, Connor Brackley, Katherine Bassett, Krissy Govertsen, Nicole Koczorek, Oliver Abele, Emily Casavant, Michael Kane, Zheng O'Neill, Tao Yang, Julia Day, Brent Huchuk, Runa T. Hellwig, and Marika Vellei. 2023. Ten Questions Concerning Occupant-Centric Control and Operations. *Building and Environment* 242 (Aug. 2023), 110518. https://doi.org/10.1016/j.buildenv.2023.110518

[16] Sophie Naylor, Mark Gillott, and Tom Lau. 2018. A Review of Occupant-Centric Building Control Strategies to Reduce Building Energy Use. *Renewable and Sustainable Energy Reviews* 96 (Nov. 2018), 1–10. https://doi.org/10.1016/j.rser.2018.07.019

[17] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning Structured Output Representation Using Deep Conditional Generative Models. In *Advances in Neural Information Processing Systems*, Vol. 28. Curran Associates, Inc.

[18] Andrew J. Sonta, Perry E. Simmons, and Rishee K. Jain. 2018. Understanding Building Occupant Activities at Scale: An Integrated Knowledge-Based and Data-Driven Approach. *Advanced Engineering Informatics* 37 (Aug. 2018), 1–13. https://doi.org/10.1016/j.aei.2018.04.009

[19] Zhe Wang, Tianzhen Hong, and Mary Ann Piette. 2019. Data Fusion in Predicting Internal Heat Gains for Office Buildings through a Deep Learning Approach. *Applied Energy* 240 (April 2019), 386–398. https://doi.org/10.1016/j.apenergy.2019.02.066