

Poster Abstract: The Smart Building Privacy Challenge

Tong Wu^{†*}, Murtadha Aldeer^{†*}, Tahiya Chowdhury^{*}, Amber Haynes^{*}, Fateme Nikseresht[◊], Mahsa Pahlavikhah Varnosfaderani[◊], Jiechao Gao[◊], Arsalan Heydarian[◊], Brad Campbell[◊], Jorge Ortiz^{*}

[†] The first two authors contributed equally to this work.

^{*}Rutgers University, [◊]University of Virginia

{tong.wu96,murtadha.aldeer,tahiya.chowdhury,amber.haynes,jorge.ortiz}@rutgers.edu,
{fn5an,mp3wp,jg5ycn,ah6rx,bradjc}@virginia.edu

ABSTRACT

Time-series data gathered from smart spaces hide user's personal information that may arise privacy concerns. However, these data are needed to enable desired services. In this paper, we propose a privacy preserving framework based on Generative Adversarial Networks (GAN) that supports sensor-based applications while preserving the user identity. Experiments with two datasets show that the proposed model can reduce the inference of the user's identity while inferring the occupancy with a high level of accuracy.

CCS CONCEPTS

- Computer systems organization → Sensor networks; Embedded systems.

KEYWORDS

Privacy, IoT, Smart building, Occupancy detection

1 INTRODUCTION

Shared, private spaces, as found in workplaces and co-working spaces are increasingly becoming instrumented [4]. Internet-of-things (IoT) devices, such as smart thermostats and inexpensive commercial and off-the-shelf (COTS) sensors – such as Doppler, temperature and humidity sensors – have been used to infer various kinds of indoor activities to improve building energy efficiency [2]. For example, Beltran et al. [3] uses a new sensor that fuses a low-cost sensors to estimate occupancy and set the HVAC control, achieving savings up to 25% annually. Agarwal et al. [1] use a mix of sensors to infer occupancy and perform prediction over a time horizon, which is also used for intelligent control and save between 10%–15% on average for various types of buildings.

While broad sensing in spaces can be used to improve efficiency and provide new capabilities for space and HVAC activation management, they can also reveal highly sensitive information about occupants. Indeed, numerous studies have shown that highly personal information can be inferred from data with seemingly little

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

BuildSys '21, November 17–18, 2021, Coimbra, Portugal

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9114-6/21/11.

<https://doi.org/10.1145/3486611.3492234>




Figure 1: A high level architecture for our privacy preserving model.

interpretable content [7]. Certainly, the choice has been overwhelmingly in favor of broadening the use of sensors in building and the trend is set to accelerate. Privacy remains a major concern for future building with deeply sensed environment. However, we conjecture that the inherent tension between privacy and utility in smart buildings can be addressed through machine learning techniques and should be explored further to examine the broader implications and theoretical limits of this trade-off. In this paper, we present a GAN-based [6] framework for supporting sensor-based applications while preserving user privacy. We pick our application class to be a (set of) classifier(s) that are used to infer the general state of a space, such as occupancy detection or count. We call this the *operational classifier*. We also consider privacy exposing process to be represented by a classifier that is able to reveal personal information, such as who is in the room. We call this the *expositional classifier*, since it exposes personal information.

2 APPROACH

We propose a GAN-based framework (Figure 1) that assumes the operational and expositional classifiers have been trained on the original data. It learns to mimic the input data but add noise to it such that the performance of the operational classifiers are maintained but the expositional ones are not. Given data collected by IoT devices, D , we consider two applications that can be enabled: occupancy detection and occupant identification. To support these applications, two classifiers can be trained (operational classifier and expositional classifier), C_1 and C_2 . Assuming these classifiers provide given performance levels on the original time-series data, $P(C_1, D) = A_1$ and $P(C_2, D) = A_2$, respectively; we aim to modify the data by concealing the user's sensitive portions, such that $P(C_1, \tilde{D}) \geq P(C_1, D)$, while the performance of the expositional classifier is $P(C_2, \tilde{D}) < P(C_2, D)$.

We designed a GAN [6] to generate synthetic data (Figure 2). In our network, the generator takes a random noise vector as input and generates fake data, which has the same dimension with real data. Then, it is send to discriminator and compared with real data.

Table 1: Performance of the occupancy classifier.

Dataset	Original data		Synthetic data	
	Acc.	F1-score	Acc.	F1-score
A	0.91	0.934	0.887	0.917
B	0.927	0.85	0.892	0.883

In the meantime, the generated data is classified by two application classifiers we mentioned before, given the labels. Two classifiers output 2 different losses. We added these losses into generator. We utilized a customized loss combination function to guide generator by reducing loss from GAN and operational classifier (C_1) and increasing loss from expositional classifier (C_2). It eventually tries to optimize the C_1 and break down the C_2 .

3 EVALUATION AND PRELIMINARY RESULTS



We adopted Support Vector Machines (SVM) with linear kernels for two types of classifiers. We developed the generator and discriminator networks by fully connected layer with 3 hidden layers. Each layer has 64 units and is activated by ReLU function.

We collect two datasets using different sensor devices in different universities. The first dataset is collected using Maestro [5], a custom designed prototype that embeds 9 sensors and powered by a Raspberry Pi3B+[®]. The sampling rate for the sensors is 30 samples/sec. There are three subjects in this dataset that have preformed a set of activities in an office setting, individually (typing on computer, calling via Zoom[®], writing on a white board, and drinking while standing). Three sensing units were deployed in an office and each subject was present for three sessions that are 1 hour long, where each activity lasted for 15 minutes on average. That resulted in 9 hours of data collected over different days. Data labels were provided by the subjects.

The second dataset is collected using AWAIR Omni[®] which collects indoor environmental quality (IEQ) factors data (e.g. CO₂, PM2.5, illuminance, etc.). 5 subjects conducted this set of experiments, where they performed the same actions as dataset A with minor differences. The experiments were run twice with each participant, each run taking 1 hour with sensors providing samples every 10 seconds. Each activity was performed for 15 minutes.

In our evaluation, we used Principal component analysis (PCA) and t-Distributed Stochastic Neighbor Embedding (t-SNE) methods to reduce the high dimensions in the data and visualizing the classes relationship in the data. Hence, most of the synthetic data points were similar to the original ones. In comparison, synthetic samples generated by GAN-based framework appeared in only one cluster of PCA. This might be the crucial feature that GAN maintained for operational classifier to identify whether the space is occupied.

The performance of the operational classifier is depicted in Table 1. Although the performances of most of the metrics drop for

**Figure 2: Architecture of our GAN.****Figure 3: Confusion matrices for the expositional classifier using (a) & (c) the original data from A & B dataset, and (b) & (d) the corresponding synthetic data generated by our GAN model. Classes refer to Null (empty office), subject 1, subject 2, etc, respectively.**

around 2% with synthetic data, it shows our model is able to achieve comparable accuracies and F1-scores with the original data and the synthetic data. Figure 3 shows the performance of the expositional classifier relate to both datasets. The synthetic data generated by our model misleads the expositional classifier to consider every occupant as the same one (occupant #1) on dataset A. For dataset B, the classifier misclassified occupant #1, #2, #3, #5 but identified #4 correctly with a slightly higher prediction score (0.54 vs 0.46). It might indicate that occupant #4 left intense distinctive characteristics in the tested space and they are stubborn to remove. In general, these results show that the proposed model can anonymize the occupants identity efficiently while detecting the occupancy within the smart space.

REFERENCES

- 1] Yuvraj Agarwal et. al. 2010. Occupancy-driven Energy Management for Smart Building Automation. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building (BuildSys '10)*, 1–6.
- 2] Bharathan Balaji et. al. 2013. Sentinel: Occupancy Based HVAC Actuation Using Existing WiFi Infrastructure Within Commercial Buildings. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems (SenSys '13)*, Article 17, 14 pages.
- 3] Alex Beltran and Alberto E. Cerpa. 2014. Optimal HVAC Building Control with Occupancy Prediction. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings (BuildSys '14)*, 168–171.
- 4] BuiltSpace. 2021. BuiltSpace and Buildings IOT team up to help WeWork reduce the risk of COVID-19 in buildings. <https://www1.builtspace.com/2020/12/02/builtspace-and-buildingsiot-team-up-to-help-wework-reduce-the-risk-of-covid-in-buildings/>.
- 5] Tahiry Chowdhury and Murtadha Aldeer et. al. 2021. Poster: Maestro—An Ambient Sensing Platform With Active Learning to Enable Smart Applications. In *Proceedings of the 2021 International Conference on Embedded Wireless Systems and Networks (EWSN '21)*, Article 15.
- 6] Ian Goodfellow and et. al. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- 7] Jacob Kröger. 2019. Unexpected Inferences from Sensor Data: A Hidden Privacy Threat in the Internet of Things. In *Internet of Things. Information Processing in an Increasingly Connected World*, Leon Strous and Vinton G. Cerf (Eds.). Springer International Publishing, Cham, 147–159.