CIPhy: Causal Intervention with Physical Confounder from IoT Sensor Data for Robust Occupant Information Inference

Zhizhang Hu
zhu42@ucmerced.edu
University of California, Merced
California, USA

Tong Yu
worktongyu@gmail.com
Carnegie Mellon University
Pennsylvania, USA

Ruiyi Zhang
ryzhang.cs@gmail.com
Duke University
North Carolina, USA

Shijia Pan
span24@ucmerced.edu
University of California, Merced
California, USA

ABSTRACT

Occupant information inference with IoT sensor data enables many smart applications, such as patients'/older adults' in-home monitoring. The difficulty of collecting labeled real-world IoT sensor data often leads to reliability and scalability issues for those systems. Extensive prior works (e.g., domain adaptation) focus on the domain shift issues, i.e., the inconsistent data feature and label relationship, and dataset bias is often neglected. Dataset bias is commonly caused by limited and varied accessibility to labeled data for each class, and it is inevitable for real-world datasets. The model trained with a biased dataset fits into the bias, hence cannot further generalize to the testing data for accurate inference.

We propose *CIPhy*, a causal intervention scheme with physical confounders measured from the sensor data to achieve robust occupant information inference. We model the dataset bias as a confounding problem. There exists a confounder directly impacts both data feature and label, and each class's accessibility to labeled data varies when the confounder's condition changes. The model trained with biased data learns a spurious feature-label correlation conditioned on the confounder's condition in the training data. When testing data has a different condition, i.e., confounding shift, this correlation can not be applied. By using the causal intervention, e.g., backdoor adjustment, the confounding shift's negative impact on the data-driven models can be mitigated. The *CIPhy* decouples the sensor data to measure the confounder, then conducts the causal intervention for a de-biased occupant information inference. We use a public dataset on occupant identification as a case study, to investigate the feasibility of applying causal intervention to resolve the dataset bias issue. From the experiment, *CIPhy* achieves up to 11.42% identification accuracy improvement compared to baselines given the biased training data and confounding shift.

CCS CONCEPTS

• Human-centered computing → Ubiquitous and mobile computing; • Computing methodologies → Causal reasoning and diagnostics.

KEYWORDS

Occupant Information Inference, Causal Intervention, Dataset Bias

ACM Reference Format:
Zhizhang Hu, Tong Yu, Ruiyi Zhang, and Shijia Pan. 2022. *CIPhy*: Causal Intervention with Physical Confounder from IoT Sensor Data for Robust Occupant Information Inference. In *The 20th ACM Conference on Embedded Networked Sensor Systems (SenSys '22), November 6–9, 2022, Boston, MA, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3560905.3568304

1 INTRODUCTION

The development of Internet-of-Things (IoT) systems enables occupant information inference for various smart applications, such as personalized HVAC/light control, and patient/older adults' in-home monitoring [1]. A variety of sensing systems have been explored, including but not limited to vision-, audio-, vibration-, and wearable-based approaches [5, 6, 12, 15, 16]. Most of these systems rely on data-driven machine learning models trained with labeled data to conduct occupant information inference [23]. However, in real-world scenarios, pure data-driven approaches are not scalable because they are not robust to the challenges raised by dataset properties (e.g., domain shift, dataset bias). Extensive priors studies have investigated **domain shift** issues, which focus on the relationship inconsistency between the data feature and label. On the other hand, **dataset bias**-caused scalability issue is often neglected. Dataset bias commonly results from limited and imbalanced training data and/or labels. For example, under certain physical conditions, data/labels from one class may have different accessibility than another. A model trained with biased data may learn a spurious correlation between the label and the physical condition, and relays on this correlation for inference [11, 24]. Prior study has shown that ignoring such dataset confounding can introduce omitted variable bias [13], thus domain adaptation algorithms that focus only on the correlation between data feature and labels are inapplicable for resolving dataset bias. We adopt **causal intervention** from causal inference [18] to mitigate the dataset bias' impact on learning accuracy, which is referred to as the **confounding shift** [11], and to learn an unbiased correlation [18].

We present *CIPhy*, a Causal Intervention scheme with Physically confounders measured from the sensor data for IoT systems. We first identify major confounding physical factors for the given datasets. Then, *CIPhy* decouples and measures the confounder from the IoT sensor data. Next, *CIPhy* leverages the causal intervention – backdoor adjustment – to achieve a de-biased occupant information inference. We use a public occupant identification dataset with dataset bias conditions to evaluate our proposed scheme. Our scheme achieves up to 11.42% identification accuracy improvement

SenSys '22, November 6–9, 2022, Boston, MA, USA

ACM ISBN 978-1-4503-9886-2/22/11...$15.00
https://doi.org/10.1145/3560905.3568304

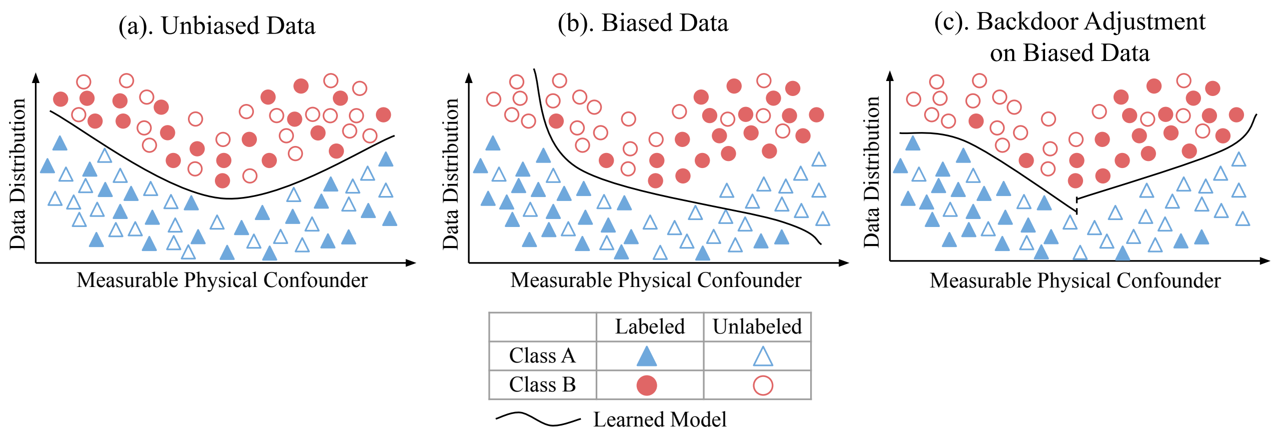

Figure 1: Examples of the dataset bias and its impact on the model. The distribution of two classes (y-axis) changes along the measurable physical confounder (x-axis) in all three sub-figures. In (a), along the x-axis, there is evenly distributed labeled data. Hence the learned model is unbiased. In (b), the labeled data is not balanced for different measurable physical confounders, thus the model biasedly fits the data. In (c), backdoor adjustment stratifies the data based on the measurable physical confounder. For each stratification, the model is adjusted to be unbiased.

over six different dataset bias conditions. We further provide a thorough verification of the impact of the causal intervention from the feature dependency perspective. To the best of our knowledge, this is the first work that uses causal intervention for robust occupant information inference. The key contributions of our proposed schemes are as follows.

- We demonstrate that confounding shift in the biased IoT dataset causes scalability issues of occupant information inference models.
- We propose *CIPhy*, a scheme that decouples the confounder from the sensor data. By measuring the confounder, *CIPhy* conducts causal intervention on the inference model to learn an unbiased data-label correlation.
- We conduct a preliminary feasibility check on a real-world public dataset to verify the effectiveness of the causal intervention, and explain the mechanism of the intervention.

2 BACKGROUND: CONFOUNDING

Conventionally, data-driven occupant information inference models are trained to learn the correlation between the observed sensor signal pattern (X) and the according label (Y). The inference is then based on the probability $P(Y|X)$. However, for real-world IoT systems, this correlation is conditional when both X and Y are influenced by another variable Z [18]. In causal inference, this is referred to as confounding, and the variable Z is the confounder. Figure 1 (a) depicts the confounding and the conditional X-Y correlation. Along the change of the measurable physical confounder (x-axis in Figure 1 (a)), the data distribution of two classes (y-axis) also changes. To train a generalize-able inference model, this conditional correlation requires the labeled training data represents the overall data distribution for each $z \in Z$. In Figure 1 (a), the labeled data evenly distributes along the measurable physical confounder axis, hence the learned model can generalize well to the unseen testing data. However, the labeled data is usually limited and have varied accessibility for each class in a real-world system. If the labeled data does not evenly contain confounder's each condition, there is a dataset bias, illustrated in Figure 1 (b). In the figure, most of class A's labeled data lies towards the left side of the measurable physical confounder axis, while a tiny amount of labeled class A data lies on the right side. On the other hand, class B's labels are majorly towards the right side of the measurable physical confounder axis, and a limited amount of labeled class B data lies on the left side. The model trained with biased training data fits into the bias in each $z \in Z$, and learns a relationship between the label Y and confounder Z for inference [24]. However, this relationship is built on data bias and is not generalizable [17]. When the model is applied to testing data with a different bias than the training data, the learned relationship negatively impacts the inference accuracy. This difference in the bias is named as confounding shift and described as $P_{train}(Y|Z) \neq P_{test}(Y|Z)$ [11]. Note that confounding shift is not the same issue as domain shift. Domain shift is the data feature X and label Y are drawn from different joint distributions, i.e., $P_{train}(X, Y) \neq P_{test}(X, Y)$. It can be illustrated as using left side of the data in 1 (a) to train the model and test on the right side of the data. Domain shift results from the distribution change, while the confounding shift results from the confounding and dataset bias.

3 *CIPHY* DESIGN

We present *CIPhy*, a causal intervention scheme that measures physical confounders from the IoT sensor data for robust occupant information inference. First, the sensor acquires the signals caused by occupants. Then *CIPhy* analyzes the signal by 1) measuring the confounder, and 2) establishing occupant information inference models. We select the confounder that is directly measurable from IoT sensor signals without the inference model. Note that a real-world dataset may contain in-definite numbers of confounders, they can be either explicit or latent. Confounder discovery is a challenging research question, which is out of the scope of this study. Even though there is not a statistical test for verifying a confounder [17], we can conceptually justify this confounder with Independent Causal Mechanisms (ICM) principle [20]. The ICM principle assumes the generative process of the cause is independent of the causal mechanism (the effect): they do not share information and each one can be manipulated independently of each other [20]. However, this kind of independence is generally violated in the opposite direction [9]. Therefore, we can leverage this property to justify the causal relationship among data features, data labels and the proposed confounder. In addition, ICM renders the system a theoretic base for focusing on one confounder and its descendants to estimate the bias resulting from this confounder [20].

CIPhy decouples the sensor data to measure the confounder Z. The measurement enables the *CIPhy* scheme to estimate the extent of the bias in the dataset and the implementation of the backdoor adjustment. After the measurement, the *CIPhy* scheme then discretizes the result into multiple levels. The discretization resolution depends on the nature of Z, i.e., if Z has a pervasively acknowledged discretization standard (e.g., a person's age can be represented as integers), we can directly adopt it. On the other hand, domain knowledge is needed to discretize the observed confounder. This discretization enables the system to decompose the correlation of $P(Y|X)$ conditioned on the confounder Z with the Bayes' rule in to the following equation:

$$P(Y|X) = \sum_{z} P(Y|X, z)P(z|X), \tag{1}$$

where $z \in Z$ is an arbitrary discretized level of measured Z. For each level of z, there is a probability $P(Y|X, z)$ for inferring the identity conditioned on signal and confounder. And the final inference is based on summing over results from all $z \in Z$, with a weight given by $P(z|X)$. In a biased dataset, there is a dominated level of z, and the model is inclined to predict a bigger $P(z|X)$ than other minority levels in Z [24]. Therefore, more weight will be given to $P(Y|X, z)$ from the dominated z level. This decomposition reveals the confounder Z introduces the bias to the inference via $P(z|X)$ and provides the capability to conduct the causal intervention.

The backdoor adjustment is then used to mitigate the impact of the confounder on the models. Backdoor adjustment is a causal intervention method that can cut off the causal relationship between X and Z, when a set of variables Z satisfies the backdoor criterion relative to X and Y [18]. Intrinsically, the backdoor criterion makes sure there are no non-causal paths between X and Y in the backdoor adjustment process. Therefore, backdoor adjustment enables the *CIPhy* to model the correlation between X on Y without the influence of Z [10]. Figure 1 (c) illustrates how the backdoor adjustment ameliorates the dataset bias' impact by stratifying the data for different measurable physical confounders and adjusting the model to be unbiased in each stratification. The backdoor adjustment is expressed as follows:

$$P(Y|do(X)) = \sum_{z} P(Y|X, z)P(z), \tag{2}$$

where $do(X)$ denotes the do-calculus [18], indicates the intervention on X. Compared to Eq. 1, in Eq. 2 z is no longer affected by X. Therefore the intervention pushes the inference to incorporate every z fairly, subject to its prior $P(z)$.

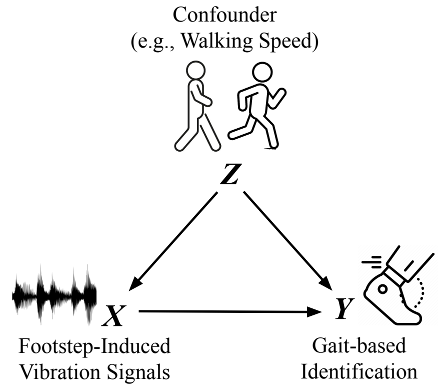

Figure 2: The directed causal model for the modeling process of footstep-based occupant identification. A confounder variable Z causes and impacts both sensor signal pattern X and occupant identity label Y.

We use a public dataset on **occupant identification as a case study** to explain the *CIPhy* design in the following section.

4 EVALUATION AND ANALYSIS

To demonstrate the effectiveness of the causal intervention on real-world IoT sensor data, we use an occupant identification dataset named FootprintID [4, 16] as a case study. Figure 2 depicts a directed causal graph of the footstep-based identification process via floor vibration. The X is the footstep-induced floor vibration signal pattern and Y is the gait-based occupant identity label. The $X \rightarrow Y$ stands for the model inferred by $P(Y|X)$. The walking speed is the confounder Z – when it changes, both vibration signal pattern X and gait pattern Y will change. We apply the ICM analyzing framework [9] to verify the confounding relationship. Taking $Z \rightarrow Y$ as an example, the generating process of Z is independent of the generating process of Y. In our occupant identification task, the measurement of Z (time interval between two detected footstep events) is independent of the measurement of Y (gait-based identity). Therefore, this causal relationship follows the ICM principle.

The dataset consists of the data from a vibration sensor placed on the floor, where 10 participants pass by the sensor with seven levels of walking speeds 10 times. The speed class 1 indicates the shortest step interval (i.e., fastest walking speed) and the speed class 7 means the longest step interval. Prior work using this dataset builds on the assumption of each person has the same amount of labeled data, which is not practical. Therefore, we prepare our training and testing datasets with confounding shift settings (Section 4.1). Then we introduce the implementation of the *CIPhy* to train and test on the aforementioned datasets (Section 4.2). Next, we present the learning results (Section 4.3). Finally, we verify the effectiveness of the causal intervention with the analysis of model's dependency on the walking speed-related features (Section 4.4).

4.1 Confounding Setting

In order to simulate the confounding shift with the public dataset, we split the training and testing data to have a different $P(Y|Z)$. In each **trial**, we select two speed classes' data, one of them is from the neutral speed, denoted as D_4 and the data from the other speed class (walk slower or faster than neutral) denoted as D_k. We randomly select five persons out of a total of 10, named as the Group 1, and the rest five persons are referred to as Group 2. Within Group 1's five person, we random sample $n\%$ ($n \in [85, 95]$) data from D_4 and $(100 - n)\%$ data from D_k, as a part of each data's training split. For Group 2's five person, we random sample $(100 - n)\%$ data from D_4 and $n\%$ data from D_k, to form the training data together with Group 1's training split. All remaining data from D_4 and D_k are combined as the testing data. With this kind of split, $P(Y)$ and $P(Z)$ are consistent in training and testing data, we can focus on varying $P(Y|Z)$ and investigate its effect. Also, for both D_4 and D_k's train split, the class is biased: there are five persons' class has nine times more data samples than the other five. On the other hand, the class is balanced for the testing data. We use accuracy as the evaluation metric. In each round, a $n \in [85, 95]$ is randomly generated for each class. We report the average value of 600 trials' accuracy values, i.e., 100 trials per speed pair.

4.2 *CIPhy* Implementation

In each trial, we first train and tune a probabilistic classification model on the training data. Following the backdoor adjustment implementation protocol [11], we incorporate a one-hot encoded feature of speed class (4 or k) to explicitly express the joint condition $P(Y|X, z)$. Then in the inference step, we follow Eq. 2 to predict the occupant identification on the test data. We estimate $P(Y|X, z = 4)$ and $P(Y|X, z = k)$ separately by setting the one-hot encoded feature as the speed class's encoding for the entire test data. Since we can measure the footstep interval, we can compare the measured intervals between the testing data and training data, then assign a speed class label to each test data sample. In this way, we can utilize the Maximum Likelihood Estimation to estimate the $P(z)$:

$$P(z = k) = \frac{\sum_{i \in |D_{Test}|} \mathbf{1}[d_i^z = k]}{|D_{Test}|}, \tag{3}$$

where $\mathbf{1}[]$ is an indicator function, $|D_{Test}|$ is the size of test data, and d_i^z is the speed class label of ith data sample in the test data. Following Eq. 2, we calculate the prediction probability $P(Y|do(X))$ for each class. The class with the highest probability is the final predicted identity.

We adopt a shallow machine learning model as the base model in this paper for its capability of investigating the model's dependency on different features for further investigating causal intervention mechanisms. Unlike the FootprintID paper[16], we adopt the Logistic Regression (LR) model as the *CIPhy*'s base model for its probabilistic property. Its probabilistic nature enables estimating the prediction likelihood. Please be notified that selecting any specific machine learning model and data feature is not our focus. We use the LR as one exemplary machine learning model, and *CIPhy* can be smoothly extended to other machine learning and deep learning models. To evaluate the effectiveness of *CIPhy*, we set two baselines for comparison: *1) Logistic Regression*: For a fair comparison, we select *CIPhy*'s base model as a baseline. We consider **the accuracy difference between this baseline and *CIPhy* reveals the negative impact of the confounding shift**. *2) Correlation Alignment* (CORAL): CORAL is an unsupervised domain adaptation algorithm [22]. CORAL aligns the second-order statistics of feature distributions between two datasets, without the need for labels. We select CORAL for its simplicity and proven effectiveness [22].

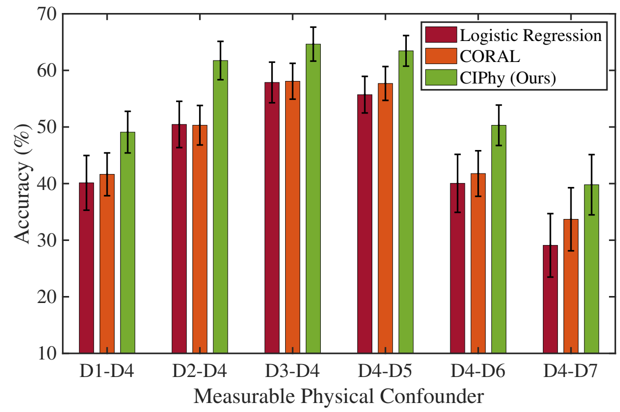

Figure 3: Occupant identification accuracy for dataset with different measurable physical confounder settings (i.e., different walking speed pairs). *CIPhy* achieves the highest accuracy in all investigated settings, indicating its robustness against confounding shift.

4.3 Result Analysis

Figure 3 shows the occupant identification accuracy averaged over 600 trials with different D_4-D_k pairs. We observe that *CIPhy* consistently outperforms two baseline models on all six pairs. It verifies that the domain adaptation designed for the domain shift is not efficient for resolving the confounding shift. When two speeds are close, e.g., D_4 and D_3, all three models can achieve a relatively high accuracy: 57.87%±3.59% for Logistic Regression, 58.08%±3.16% for CORAL, and 64.64%±3.00% for *CIPhy*. When the speed difference increases, the two baseline models show more significant decreases compared to *CIPhy*. Take the pair of D_4 and D_2 as an example, the Logistic Regression achieves an accuracy of 50.45 %±4.08% (7.42% decrease), the CORAL has an accuracy of 50.31%±3.48% (7.77% decrease), and the *CIPhy*'s accuracy is 61.74%±3.38% (2.90% decrease). While the walking speed is significantly different in the pair, e.g, D_4 and D_7, the Logistic regression model's accuracy decreases drastically to 29.08%±5.60%. The CORAL model outperforms the Logistic Regression and achieves an accuracy of 33.67%±5.73%. Our *CIPhy* has the least negative impact from the dataset bias and achieves the highest accuracy of 39.79%±5.32%.

From a causal perspective, when the model is trained on biased data, the model learns a spurious relationship between the walking speed feature and the identity. For example, if in the training data person A has a lot of data from speed 4 while a tiny amount of speed k data, the model tends to relate speed 4 with identity A. During the model's inferences on the testing data, the decision is biasedly impacted by that speed-identity-relationship. When the feature's distribution changes drastically, it is more challenging for the model to conduct inference with the signal feature, hence it relies more on the spurious relationship learned from training data. Therefore, baseline models' accuracy decreases drastically when

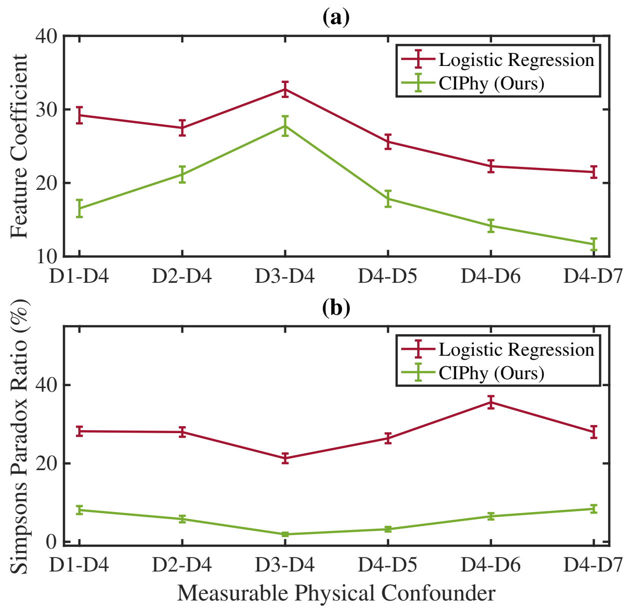

Figure 4: Backdoor Adjustment mitigates the negative impact of confounding shift. (a) shows the summation of the step frequency features' coefficient absolute value of 10 target classes. The lower the value in each pair, the less dependency on the confounder for inference. (b) shows the ratio of Simpson's Paradox for the step frequency feature coefficient when predicting 10 persons. The lower the value, the less the model is influenced by the confounder.

the walking speed difference increases. We will dive deep into the causal intervention in the following section.

4.4 Dive into Backdoor Adjustment's Intervention

To understand the mechanism of backdoor adjustment on mitigating the negative impact of confounding shift, we investigate from two perspectives: 1). logistic regression's coefficient value of walking speed-related feature, and 2). the ratio of Simpson's paradox is shown on the sign of the walking speed-related feature coefficient.

A trained logistic regression model has a coefficient(s) for each feature in the data [19]. This coefficient of each feature determines the strength of correlations between the feature and predicting target, i.e., the dependency on each feature [3, 19]. The closer one feature's coefficient value is to zero, the less influence this feature has on the model inference [14]. In Figure 4 (a), we plot the summation of the step frequency feature's coefficient absolute value of 10 target classes from two models, for all dataset pairs. Since the step frequency is directly related to the walking speed for each person, the coefficient of this feature indicates the model's dependency on the walking speed. Figure 4 (a) shows Logistic Regression model's step frequency feature coefficient is consistently higher than *CIPhy*'s. This indicates the *CIPhy* is less depending on the step frequency feature, hence the negative impact from the walking speed is mitigated [10]. Also, the gap between the two models' coefficient summation increases along with when the walking speed difference: 4.99 for the pair D_4-D_3 and 12.67 for D_4-D_1. This shows that the baseline model relies more on the spurious relationship between the walking speed and identity **when the feature's distribution changes drastically**, and is affected by the confounding shift more significantly.

Simpson's Paradox is a statistical phenomenon where an association between two variables (X, Y) emerges in a general population and disappears or reverses when the general population is divided into smaller sub-populations conditioned on a third variable Z [21]. Prior study has shown that the appearance and the extent of the Simpson's Paradox in model's parameters can be used to quantify the model's impact from the confounding shift [10]. Taking the logistic regression coefficient as an instance, the sign of the coefficient determines the direction of the relationship between X and the prediction probability of $P(Y|X)$ [19]. For an arbitrary class y, when β is greater than zero, a bigger X value is associated with a bigger prediction probability of $Y = y$. On the other hand, if β is smaller than zero, a bigger X value is associated with a smaller prediction probability of $Y = y$. Therefore, for the walking speed feature coefficient, we can compute the ratio it exhibits Simpson's paradox when predicting each person's identity. For example, when the model is trained on the entire training data, the step frequency feature coefficient for predicting identity as person A is a positive value. However, if the model is trained on D_4 or D_k's training split, the coefficient changes to be a negative value. When this phenomenon happens, we consider this coefficient shows Simpson's Paradox. Figure 4 (b) plots the ratio of experiment rounds that shows Simpson's Paradox for the walking speed feature coefficient when predicting 10 persons. From the figure, we observe that this ratio for the Logistic Regression model is substantially higher than that of *CIPhy* for all six pairs. The mean value of ratios for Logistic Regression is between 20% to 40%, while the *CIPhy*'s ratios are all less than 10%. It indicates that the backdoor adjustment effectively reduces the negative impact from the confounder and improves identification robustness with a biased dataset.

5 RELATED WORKS

5.1 Occupant Information Inference

Occupant information inference is a fundamental base for many smart applications, like personalized comfort adjustment [5] and occupant behavior analysis [12, 15]. Prior works investigate information inference with different sensing modalities. Francis et al. present a depth-based system for inferring occupants' thermal comfort [5]. The system leverages the depth image of body shape to predict occupants' thermal feeling and generates the HVAC temperate setting point for each occupant. Mollyn et al. introduce a multimodal sensing system with the IMU sensor and microphone in the smartwatch to infer the context of daily living and the occupant activity [15]. Zeng et al. present a framework that can identify an occupant's identity with WiFi Channel State Information (CSI) [25]. They first identify a person's walking gait from the CSI pattern, then further infer the person's identity. These prior works conduct the information inference based on the association between the feature and the label, which is vulnerable to bias in the dataset.

5.2 Causal Learning

Causal learning, sometimes referred to as causal machine learning, is an intersection of causal inference and machine learning (ML) [2]. Common ML algorithms learn the correlation-based patterns and relationships from data for inference [24]. However, data may contain spurious correlations hence negatively affecting ML algorithms' performance. With tools like causal interventions, causal learning brings opportunities to ML research for reducing the influence of bias during inference has achieved promising results in common ML domains, such as Computer Vision [24, 26] and Natural Language Processing [7, 10]. Wang et. al. introduces an unsupervised feature representation learning method, which leverages backdoor adjustment to reduce the bias of the image background on the object representation learning [24]. Zhu et. al. proposes to leverage the backdoor adjustment to mitigate the bias of class imbalance and train a robust unbiased model for long-tailed classification tasks [26]. Garg et. al. utilizes counterfactual intervention to reduce the bias from sensitive attributes in the text data and build a fair text classifier [7]. Given the fact that causal learning is well studied in these domains, it is still in an emerging phase in the IoT domain. Hu et. al. peeks at the dataset bias and fairness in the infrastructural sensing data, however, they do not propose methods to tackle the bias in the dataset [8]. In this paper, we show the feasibility to leverage causal learning to build a robust ML model for the IoT domain's sensory data.

6 DISCUSSIONS AND FUTURE DIRECTIONS

In this paper, we aim to demonstrate the feasibility of using causal intervention to enhance occupant information learning given the sensor data bias in the built environment. We further discuss the future directions in this section.

6.1 Confounder Discovery

The example confounders explored in this work are selected with domain knowledge of structural vibration and human behavior. In real-world IoT sensing systems, various factors may result in a different cause-effect relationship. Thus, the types of bias induced by the confounders can be also various. Also, there may be causal relationships between confounders in the dataset, which makes the problem more difficult. One future direction is how can we identify the confounder and extract the causal relationship between variables in the data. One potential solution is infusing the physical knowledge of the sensing system into current statistical causal discovery methods, and letting the real-world knowledge drive the process of confounder identification. Additionally, the latent representation encoded by deep learning models raises another challenge: there may be bias and confounders in the latent representations that cannot be understood by humans. We can explore the use of the first laws of physics as a way to automate the discovery of confounder variables that may be prevalent in IoT applications.

6.2 Multimodal Causal Learning

In this paper, we focus on causal learning with a single sensing modality. Recent studies have shown that multiple sensing modalities can provide complementary information and provides higher performance than uni-model systems. On the other hand, applying causal learning to multimodal sensing systems is more challenging than that to uni-modal systems. The first challenge comes from: how can we model the system with a unified mathematical description? The signal semantics can varies significantly in a multimodal system. For example, Electromyography (EMG) sensor measures the electrical activity of muscles, while the infrastructural sensors measure a person's interaction with the environment. The second challenge comes from the causal relationship between each sensor's target. For example, EMG can measure the muscle activity of one person's legs, those legs' activities can cause footstep vibration sensed by the vibration sensor. How to decouple the cause-effect between sensing modalities is a question to be answered. The third challenge comes from whether the confounders can be measured for all sensing modalities. Unlike footstep intervals, some variables can be difficult to be measured, like the fatigues and stress of the muscle. How can we estimate those confounder properties is a challenge for applying causal learning to multimodal systems.

7 CONCLUSION

In this paper, we investigate the feasibility of applying causal learning to IoT-based occupant information inference problems. We propose *CIPhy*, a causal intervention scheme with measurable physical confounders from the sensor data to achieve robust occupant information inference with dataset bias. We model the dataset bias as an issue raised by confounding. The model trained with biased data learns a spurious feature-label correlation conditioned on the confounder's condition in the training data, which can not be generalized to the testing data. We leverage the backdoor adjustment to conduct a causal intervention on the model for mitigating the negative impact of the bias. From the evaluation with a real-world public dataset for occupant identification, the proposed causal intervention approach outperforms baseline models consistently. In the quantitive analysis of the model, we find the proposed causal intervention scheme effectively lowers the model's dependency on the spurious relationship between the confounder and the data label, and increases the identification robustness.

REFERENCES

[1] Bharathan Balaji, Jian Xu, Anthony Nwokafor, Rajesh Gupta, and Yuvraj Agarwal. 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*. 1–14.

[2] Huigang Chen, Totte Harinen, Jeong-Yoon Lee, Mike Yung, and Zhenyu Zhao. 2020. Causalml: Python package for causal machine learning. *arXiv preprint arXiv:2002.11631* (2020).

[3] Paul D Cleary and Ronald Angel. 1984. The analysis of relationships involving dichotomous dependent variables. *Journal of Health and Social Behavior* (1984), 334–348.

[4] Yiwen Dong, Shijia Pan, Tong Yu, Mostafa Mirshekari, Jonathon Fagert, Amelie Bonde, Ole J. Mengshoel, Pei Zhang, and Hae Young Noh. 2021. *The FootprintID Dataset: Footstep-Induced Structural Vibration Data for Indoor Person Identification with Different Walking Speeds.* https://doi.org/10.5281/zenodo.4691144

[5] Jonathan Francis, Matias Quintana, Nadine Von Frankenberg, Sirajum Munir, and Mario Bergés. 2019. Occutherm: Occupant thermal comfort inference using body shape information. In *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. 81–90.

[6] Davrondzhon Gafurov, Einar Snekkenes, and Patrick Bours. 2007. Gait authentication and identification using wearable accelerometer sensor. In *2007 IEEE workshop on automatic identification advanced technologies*. IEEE, 220–225.

[7] Sahaj Garg, Vincent Perot, Nicole Limtiaco, Ankur Taly, Ed H Chi, and Alex Beutel. 2019. Counterfactual fairness in text classification through robustness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 219–226.

[8] Zhizhang Hu, Yue Zhang, and Shijia Pan. 2021. Footstep-Induced Floor Vibration Dataset: Reusability and Transferability Analysis. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*. 546–551.

[9] Zhijing Jin, Julius von Kügelgen, Jingwei Ni, Tejas Vaidhya, Ayush Kaushal, Mrinmaya Sachan, and Bernhard Schölkopf. 2021. Causal direction of data collection matters: Implications of causal and anticausal learning for NLP. *arXiv preprint arXiv:2110.03618* (2021).

[10] Virgile Landeiro and Aron Culotta. 2016. Robust text classification in the presence of confounding bias. In *Thirtieth AAAI Conference on Artificial Intelligence*.

[11] Virgile Landeiro and Aron Culotta. 2018. Robust text classification under confounding shift. *Journal of Artificial Intelligence Research* 63 (2018), 391–419.

[12] Gierad Laput, Karan Ahuja, Mayank Goel, and Chris Harrison. 2018. Ubicoustics: Plug-and-play acoustic activity recognition. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. 213–224.

[13] Lung-Fei Lee. 1982. Specification error in multinomial logit models: Analysis of the omitted variable bias. *Journal of Econometrics* 20, 2 (1982), 197–209.

[14] Scott Menard. 2002. *Applied logistic regression analysis*. Number 106. Sage.

[15] Vimal Mollyn, Karan Ahuja, Dhruv Verma, Chris Harrison, and Mayank Goel. 2022. SAMoSA: Sensing Activities with Motion and Subsampled Audio. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (2022), 1–19.

[16] Shijia Pan, Tong Yu, Mostafa Mirshekari, Jonathon Fagert, Amelie Bonde, Ole J Mengshoel, Hae Young Noh, and Pei Zhang. 2017. Footprintid: Indoor pedestrian identification through ambient structural vibration sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 1–31.

[17] Judea Pearl. 1998. Why there is no statistical test for confounding, why many think there is, and why they are almost right. (1998).

[18] Judea Pearl. 2009. *Causality*. Cambridge university press.

[19] Chao-Ying Joanne Peng, Kuk Lida Lee, and Gary M Ingersoll. 2002. An introduction to logistic regression analysis and reporting. *The journal of educational research* 96, 1 (2002), 3–14.

[20] Bernhard Schölkopf. 2022. Causality for machine learning. In *Probabilistic and Causal Inference: The Works of Judea Pearl*. 765–804.

[21] Edward H Simpson. 1951. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society: Series B (Methodological)* 13, 2 (1951), 238–241.

[22] Baochen Sun, Jiashi Feng, and Kate Saenko. 2017. Correlation alignment for unsupervised domain adaptation. In *Domain Adaptation in Computer Vision Applications*. Springer, 153–171.

[23] Dipti Trivedi and Venkataramana Badarla. 2020. Occupancy detection systems for indoor environments: A survey of approaches and methods. *Indoor and Built Environment* 29, 8 (2020), 1053–1069.

[24] Tan Wang, Jianqiang Huang, Hanwang Zhang, and Qianru Sun. 2020. Visual commonsense r-cnn. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10760–10770.

[25] Yunze Zeng, Parth H Pathak, and Prasant Mohapatra. 2016. WiWho: WiFi-based person identification in smart spaces. In *2016 15th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. IEEE, 1–12.

[26] Beier Zhu, Yulei Niu, Xian-Sheng Hua, and Hanwang Zhang. 2022. Cross-domain empirical risk minimization for unbiased long-tailed classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 3589–3597.