

TSMOBN: INTERVENTIONAL GENERALIZATION FOR UNSEEN CLIENTS IN FEDERATED LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Generalizing federated learning (FL) models to unseen clients with non-iid data is a crucial topic, yet unsolved so far. In this work, we propose to tackle this problem from a novel causal perspective. Specifically, we form a training structural causal model (SCM) to explain the challenges of model generalization in a distributed learning paradigm. Based on this, we present a simple yet effective method using test-specific and momentum tracked batch normalization (TsmoBN) to generalize FL models to testing clients. We give a causal analysis by formulating another testing SCM and demonstrate that the key factor in TsmoBN is the test-specific statistics (i.e., mean and variance) of features. Such statistics can be seen as a surrogate variable for causal intervention. In addition, by considering generalization bounds in FL, we show that our TsmoBN method can reduce divergence between training and testing feature distributions, which achieves a lower generalization gap than standard model testing. Our extensive experimental evaluations demonstrate significant improvements for unseen client generalization on three datasets with various types of feature distributions and numbers of clients. It is worth noting that our proposed approach can be flexibly applied to different state-of-the-art federated learning algorithms and is orthogonal to existing domain generalization methods.

1 INTRODUCTION

Tackling non-iid data is an important yet challenging problem in federated learning (FL). Previous works mainly studied how to effectively *train* the model on non-iid data (Zhao et al., 2018; Li et al., 2020; Karimireddy et al., 2019; Hsu et al., 2019; Reddi et al., 2021; Wang et al., 2020a; Li et al., 2021). It still remains under-explored regarding how to generalize the FL models to unseen non-iid data for *testing*. Here, we address this task of unseen client generalization with the feature shift (Li et al., 2021). The non-iidness induced by feature shift is critical in many real-world scenarios. For example, medical images collected from different hospitals can vary a lot in appearance due to different imaging protocols. Though there are methods addressing the related problem of domain generalization (Li et al., 2018b;a; Dou et al., 2019; Jin et al., 2020; Gulrajani & Lopez-Paz, 2020; Zhao et al., 2020), the requirement of centralizing different data distributions makes them not applicable in FL. A few domain generalization approaches could be cast to fit for FL privacy-preserving setting, however, would unavoidably affect the well-established FL paradigm and complicate the training process.

The key reason for the performance degeneration during the test phase is that a global model trained on heterogeneous feature distributions fails to be an accurate estimation for a different distribution on the unseen client. When training on non-iid data, an effective approach is to use batch normalization (BN) to normalize features into a uniform distribution (Liu et al., 2020; Andreux et al., 2020; Li et al., 2021). But during testing, BN is fixed and tends to use the estimated training statistics to normalize features. It has a risk of improperly normalizing features of unseen clients which can subsequently deteriorate the final model prediction. Therefore, adapting BN for the test client to obtain a similar uniform distribution as in training is a promising solution.

In this paper, we propose to use test-specific and momentum tracked batch normalization (TsmoBN) to solve the unseen client generalization problem for varying feature distributions. By drawing insights from causality we find that when BN layers are fixed for testing, they will improperly correlate training statistics and testing features. Such correlation is also known as the spurious relationship which is studied in causal inference. Inspired by this, we form two novel structural causal models

(SCMs) to analyze the relationships between non-iid data with feature shifts, learned features, and model predictions for FL. With the help of causal analysis, we find TsmoBN exactly serves as a causal intervention that effectively removes the spurious correlations. Specifically, TsmoBN calculates the mean and variance from features at test time for each BN layer. Meanwhile, it incorporates momentum to integrate information among different batches to obtain an accurate population-level estimation. The estimated pair of mean and variance can be interpreted as a surrogate variable for causal intervention. In addition, we theoretically show that the divergence of training and testing distributions can be reduced with our proposed TsmoBN, which yields a lower generalization error bound than using training statistics. Our main contributions are summarized as follows:

- We are the first to propose a systematic view of the SCM for FL to look into the unseen client generalization problem and conduct the causal analysis of our approach based on this SCM framework.
- We propose a simple yet effective method of using TsmoBN to mitigate non-iid features for unseen clients. We further theoretically demonstrate that our approach yields a lower bound on generalization error compared to standard testing on unseen clients, indicating a promising solution for unseen client generalization.
- We conduct extensive experiments using three public datasets with various feature distributions and client numbers. Our results demonstrate significant improvements in generalization over state-of-the-art FL methods. Last but not least, our solution is orthogonal and readily compatible with some current domain generalization methods which can be cast into most of the existing FL settings.

2 RELATED WORK

Federated learning on non-iid features: Recent FL studies have been increasingly investigating the issue of non-iid data distributions across different clients, from multiple aspects such as facilitating model optimization (Li et al., 2020; Karimireddy et al., 2019; Reddi et al., 2021), enhancing aggregation algorithm (Wang et al., 2020a; Pillutla et al., 2019), improving feature normalization (Li et al., 2021; Reiszadeh et al., 2020; Andreux et al., 2020), etc. For example, as an early work, FedProx (Li et al., 2020) extends the conventional FedAvg (McMahan et al., 2017) by adding a proximal term to help stabilize the optimization. SCAFFOLD (Karimireddy et al., 2019) suggests a new optimization algorithm to reduce variance for client-drift in its local updates. Later on, FedNova (Wang et al., 2020a) proposes to use normalized stochastic gradients to perform global model aggregation rather than the cumulative raw local gradient changes. Very recently, FedAdam (Reddi et al., 2021) introduces adaptive server optimization in FL. FedBN (Li et al., 2021) and SiloBN (Andreux et al., 2020) both propose to keep locally the BN statistics without aggregating them in the global model. However, all these existing works have focused on boosting performance for clients inside federated training, without elaborating how these FL algorithms can generalize well to unseen testing clients.

Structural causal models in deep learning: Causality, a concept originating from statistics (Wold, 1954; Eells, 1991; Marini & Singer, 1988; Greenland et al., 1999), has been actively introduced to deep learning recently. In particular, the structural causal model has demonstrated great potentials in explaining deep network properties and further inspiring new methods for solving application tasks. For example, recent generative-model based methods (Pawlowski et al., 2020; Wang et al., 2021a; Reinhold et al., 2021) use SCM with physical intervention (i.e., assigning specific values to variables) to enhance model transparency. They explicitly define image attributes as variables and use physical intervention to find cause effects between variables and images. Another way is to model critical parts for a certain task (e.g., long-tailed classification (Tang et al., 2020) and visual recognition (Wang et al., 2021b)) with SCM and figure out confounding effects that can be removed through back-door adjustment, therefore improving deep learning model performance. These successful works motivate us to consider forming an SCM view for federated learning, which has not been explored to date. Tackling this has unique challenges, because the physical intervention has a restriction that target intervention variables are directly changeable and back-door adjustment needs access to the parent variables of the targetted intervention one, which are all not feasible in the privacy-preserving FL setting. Novel solutions for causal intervention on deep neural networks have to be developed to fit for FL paradigm where the client data are distributed.

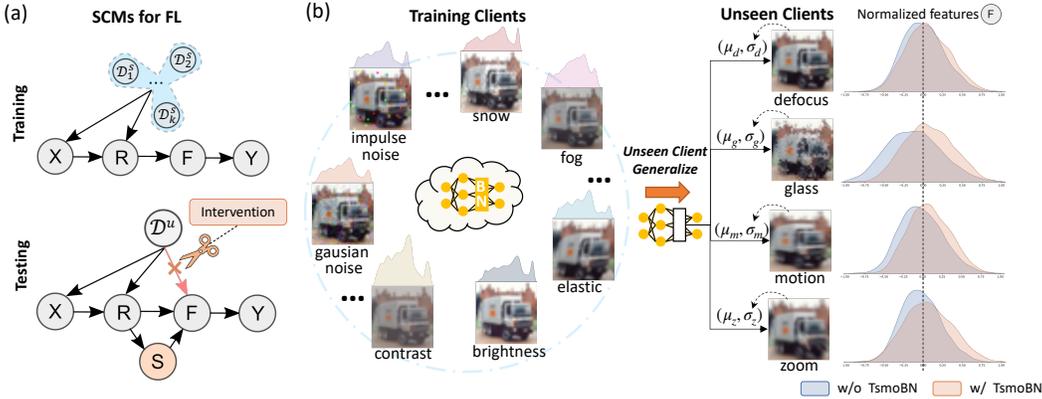


Figure 1: Overview of unseen client generalization and our proposed novel causal view. (a) Structural causal models for FL (variable notations defined in Sec. 3.1). (b) Example of training clients with various distributions and generalization on unseen clients with the illustration of change of non-iid feature distributions upon TsmoBN.

Domain generalization: Generalizing models to completely unseen domains has been an important yet challenging research topic. Domain generalization is to some extent similar to unseen client generalization, with only differences lying in centralized training or distributed training. Normalizing the learned representations to impose domain invariance has been central to domain generalization solutions. A series of methods have been proposed for aligning features across training domains to learn domain invariant representations (Jin et al., 2020; Gulrajani & Lopez-Paz, 2020), incorporating domain discriminators with adversarial learning (Li et al., 2018b; Zhao et al., 2020), and using meta-learning to simulate domain shift during training (Li et al., 2018a; Dou et al., 2019). However, these methods can not be easily fit to the FL scenario, limited by their requisite of putting multi-domain data in one place. Nevertheless, some current domain generalization methods show less restriction on distributed data, for instance, JiGen (Carlucci et al., 2019) improves generalizability by a self-supervised task of solving a jigsaw puzzle, CSD (Piratla et al., 2020) uses low-rank decomposition to extract common and domain-specific components, RandConv (Xu et al., 2021) relies on data augmentation for combating non-iid data. However, casting them to the FL setting still has to affect the well-established FL paradigm which would unavoidably complicate the training process.

3 STRUCTURAL CAUSAL MODELS FOR FEDERATED LEARNING

In this section, we start with the formulation for unseen client generalization problem (Sec 3.1) and explain the challenges by causal analysis using our formed training SCM (Sec 3.2). Then we describe our new method (i.e., TsmoBN) as a causal intervention to tackle this problem. We further analyze our method using both test SCM and generalization error bound.

3.1 PRELIMINARIES ON OUR SCMS

Notation: Let the $\mathcal{X} \subset \mathbb{R}^d$ be the input space and \mathcal{Y} denote the output space, which would be $\{-1, +1\}$ in the classification case, $\{\mathcal{D}_1^s, \dots, \mathcal{D}_K^s\}$ as the set of distributions of K distributed training clients involved in FL, and \mathcal{D}^u be the unseen testing client’s distribution¹. We consider the heterogeneous distributions that all samples X are non-iid across clients (e.g., medical images from various hospital). Specifically, for the k -th client, let $\{(x_{i,k}^s, y_{i,k}^s)\}_{i=1}^{m_k}$ be the training samples drawn independently from a training distribution \mathcal{D}_k^s with density $p_k^s(\mathbf{x}, y)$, and (x^u, y^u) be the test sample from \mathcal{D}^u with density $p^u(\mathbf{x}, y)$, decomposing the density into the marginal distribution and the conditional probability distribution, we assume there exists *feature shift* across clients as defined by Li et al. (2021). That is, we have $p(\mathbf{x})$ or $p(\mathbf{x}|y)$ varies across training and testing clients.

¹Here we present results for one unseen client \mathcal{D}^u , but the results can be easily generalized to multiples.

Problem setup: In unseen client generalization, we first obtain a global model $f : \mathcal{X} \rightarrow \mathbb{R}$ by aggregating distributed client models trained with local distributions $\{\mathcal{D}_k^s\}_{k=1}^K$. Then we apply this model f on the unseen client with the target of minimizing the empirical risk of f on samples drawn from \mathcal{D}^u . Note that the issue of generalizing to an unseen client is independent of the federated setup. Since we consider one collaboratively trained model and one unseen client at a time, we can form two SCMs analogous to centralized training and testing.

As shown in Fig. 1(a), in both SCMs we consider non-iid *inputs* X , *raw extracted features* R (i.e., the outputs of convolutional or fully-connected layers) and the *normalized features* F (i.e., the application of batch normalization to R). From F we make inferences on the *output* Y . The inputs cause the raw features which cause the normalized features which in turn cause the output. Both input and learned features are caused by the data distribution during training. We represent this relationship as directed acyclic graphs (DAGs) (Peters et al., 2017). That is, the SCMs are represented by DAGs \mathcal{G}^{tr} and \mathcal{G}^{te} , where vertices denote *variables* and directed edges denote *causal relations* (e.g., *cause* \rightarrow *effect*).

3.2 CAUSAL ANALYSIS WITHIN THE TRAINING SCM

Training SCM for FL with non-iid features: In the training phase of FL, a total of K clients aim to jointly solve the following optimization problem:

$$\min_f \left\{ \sum_{k=1}^K \frac{\alpha_k}{n_k} \sum_{i=1}^{n_k} \ell(f(\mathbf{x}_{i,k}^s), y_{i,k}^s) \right\}, \quad (1)$$

where $\alpha_k = \frac{n_k}{\sum_{k=1}^K n_k}$ is the relative sample size, and $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ denotes the loss function that measures the discrepancy between the true label and output value. As shown in Fig. 1(a), in the non-iid FL scenario, images sampled from each client local distribution ($X \leftarrow \mathcal{D}_k^s$) are assumed to be heterogeneous (e.g., the 'truck' pictures with various distributions in Fig. 1(b)). The heterogeneity may lead each client towards fitting its individual feature distributions (i.e., $\mathcal{D}_k^s \rightarrow R$) as opposed to the global objective (Karimireddy et al., 2019; Li et al., 2020), which can deteriorate the final model performance. The performance degeneration is induced by the non-iid data distribution \mathcal{D}_k^s , which is exactly known as the *confounder* that brings spurious correlations in the causality. However, since the modern neural networks typically have batch normalization layers that can normalize training data into a uniform distribution, the global model is able to converge even with the existence of the confounder. With the whole classification task defined as $X \rightarrow R \rightarrow F \rightarrow Y$ ², where the normalized features F are no longer affected by different training distributions \mathcal{D}_k^s , we finally obtain the global model that performs well on training clients. But when generalizing the model for the unseen testing client, the estimation based on training distributions may fail to be unbiased for the test distribution due to the above-mentioned feature shift, which is formulated as below:

$$\mathbb{E}_{(\mathbf{x}^u, y^u) \sim p^u(\mathbf{x}, y)} [\ell(f(\mathbf{x}^u), y^u)] \neq \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{(\mathbf{x}_k^s, y_k^s) \sim p_k^s(\mathbf{x}, y)} [\ell(f(\mathbf{x}_k^s), y_k^s)]. \quad (2)$$

Challenges: The joint distributions are different (i.e., $p^u(\mathbf{x}, y) \neq p_k^s(\mathbf{x}, y)$) due to the feature shift. To overcome the inequality between training and test estimation, many works in domain generalization and adaptation have been proposed (Arjovsky et al., 2019; Zhao et al., 2020; Mahajan et al., 2021; Zhang et al., 2020; Wang et al., 2020b), but casting these methods into FL is not trivial.

For domain generalization, a central solution is to learn domain-invariant features from multiple distributions (domains) by gathering a set of centralized non-iid datasets. Such domain-invariant feature learning can be viewed as the back-door adjustment in our training SCM. Namely, let $P(Y|do(X))$ denotes the new classifier that removes the various appearance of heterogeneous images and pursues the true causal pattern determining the label of images, this can be achieved by adjusting X given different \mathcal{D}_k^s , as shown below:

$$P(Y|do(X)) = \sum_{k=1}^K P(Y|X, \mathcal{D}_k^s) P(\mathcal{D}_k^s). \quad (3)$$

²In general, feature extraction and batch normalization will be performed repeatedly, making a long chain start from R and end at F . We instead use an arrow from R to F to simplify the graph without affecting the idea.

However, the adjustment in Eq.(3) needs to visit all the distributions to learn invariant features (i.e., sum over the product of likelihood and prior of each distribution \mathcal{D}_k^s), which is infeasible due to restriction of data sharing and local training in FL. From the back-door adjustment viewpoint, we can systematically explain why some conventional DG methods failed under the FL paradigm, and we give a formal statement on the infeasibility of applying some domain generalization methods in FL, which can be found in Appendix B.1.

For adaptation methods, directly conduct adaptation in FL training is very time-consuming. As shown in Eq.(2), to obtain the equality for a certain test distribution, we have to coordinate and adapt K clients, which suffers a high communication cost than centralized training. Hence, we argue that addressing the unseen client generalization problem at test time is a simple and less costly way that can fit the test distributions dynamically. We propose using the test-specific momentum batch normalization (TsmoBN) to effectively tackle this generalization problem and analyze our method with both test SCM and error bound.

3.3 TEST-SPECIFIC AND MOMENTUM TRACKED BN FOR UNSEEN CLIENT GENERALIZATION

Analysis with the testing SCM: When deploying the FL model that is only trained under seen distributions \mathcal{D}^s to unseen $X \sim \mathcal{D}^u$ for generalization, the difference between training and test distributions make the BN layers with the well-estimated training statistics have a risk of improperly normalizing features of unseen clients, introducing an edge $\mathcal{D}^u \rightarrow F$, as shown in Fig. 1(a). Thus, the distribution \mathcal{D}^u becomes a confounder for F that is subsequently used for the model prediction, and finally harms the test performance. To remove the confounding effects brought by \mathcal{D}^u , a direct way is using causal intervention on normalized features (i.e., $do(F)$) to let the feature distribution similar to training distributions. We achieve this intervention by introducing the surrogate variable S , which is test-specific statistics of raw features R during testing.

Test-specific and momentum tracked BN: We propose to use TsmoBN as a causal intervention to obtain the test normalized features that have similar distributions as the training normalized features. More specifically, *we calculate the mean and variance pair at test time in BN to normalize features.* The detailed algorithm is shown in Appendix D. With the help of analysis using the surrogate variable, and by adopting a popular distributional assumption that the input data points follow Gaussian distribution for all clients (Zhong et al., 2017; Li & Yuan, 2017; Ge et al., 2018; Bakshi et al., 2019; Chen et al., 2020), we can normalize features into a uniform distribution without any updates (adaptations) on trainable parameters of the obtained global model. We theoretically show that the intervention on surrogate variable S (i.e., $do(S)$) is equivalent to intervening the normalized features (i.e., $do(F)$) which can be seen as mapping the test normalized features into the same space as the training normalized features.

Since the batch-level sampling may not accurately estimate the population-level mean and variance for a client, the noticeable fluctuation in feature statistics across different batches potentially affects our assumption and lead to instability of model inference. We further propose to use momentum to integrate relations among different batches, thus reducing the variances. Precisely, given an unseen client with M batches of data to be tested in sequential (each batch has m samples and their corresponding raw features are denoted as $\{r_{1,k}^u, \dots, r_{m,k}^u\}$), we initialize $\{\mu, \sigma\}$ as $\mu = \frac{1}{m} \sum_{i=1}^m r_{i,k}$, $\sigma^2 = \frac{1}{m} \sum_{i=1}^m (r_{i,k} - \mu)^2$. With the momentum of τ , when a new batch comes in, we use the momentum as a weight coefficient gradually update the test-specific BN mean and variance. By the complete version of TsmoBN, we can approximately obtain the precise estimation for mean μ_u and variance σ_u of an unseen client. The feature normalization is shown below:

$$\begin{aligned} \mathbb{E}[\mu_u] &\approx \mu \leftarrow \tau\mu + (1 - \tau) \frac{1}{m} \sum_{i=1}^m r_{i,k}^u, & \mathbb{E}[\sigma_u^2] &\approx \sigma^2 \leftarrow \tau\sigma^2 + (1 - \tau) \frac{1}{m} \sum_{i=1}^m (r_{i,k}^u - \mu)^2 \\ TsmoBN(r_{i,k}^u) &= \frac{r_{i,k}^u - \mathbb{E}[\mu_u]}{\sqrt{\mathbb{E}[\sigma_u^2]}} \approx \frac{r_{i,k}^u - \mu}{\sqrt{\sigma^2 + \epsilon}}, & \text{for } i &= 1, \dots, m. \end{aligned} \quad (4)$$

We explain how to interpret TsmoBN as causal intervention using surrogate variable and show the equivalence of $do(S)$ and $do(F)$ in below Theorem 3.1. For vertices S, F, Y in our built DAG \mathcal{G}^{te} , S and Y are independent given F in \mathcal{G}^{te} , denoted as $(S \perp\!\!\!\perp Y | F)_{\mathcal{G}^{te}}$. Also, we denote $\mathcal{G}_{\underline{S}}^{te}$ ($\mathcal{G}_{\underline{S}}^{te}$) as the graph obtained by deleting from \mathcal{G}^{te} all arrows pointing to (emerging from) nodes in \underline{X} .

Theorem 3.1 *With the SCM defined for testing, if a variable S served as surrogate variable with following conditions satisfied: (1) F intercepts all directed paths from S to Y ; (2) $P(Y|\text{do}(S))$ is identifiable; (3) $(Y \perp\!\!\!\perp S|F)_{\mathcal{G}_{SF}^{te}}$; and (4) $(Y \perp\!\!\!\perp F|S)_{\mathcal{G}_{SE}^{te}}$, then we have*

$$P(Y|\text{do}(F)) = P(Y|\text{do}(S := (\mathbb{E}[\mu_u], \mathbb{E}[\sigma_u^2]))). \quad (5)$$

The proof of Theorem 3.1 is shown in Appendix B.2. The main idea is that the intervention on surrogate variable, i.e., $\text{do}(S := (\mathbb{E}[\mu_u], \mathbb{E}[\sigma_u^2]))$ (here $(\mathbb{E}[\mu_u], \mathbb{E}[\sigma_u^2])$ is assumed to be an unbiased estimation for the mean and variance pair of test-specific data), has the same effect on normalized feature F as the passive observation $S = (\mathbb{E}[\mu_u], \mathbb{E}[\sigma_u^2])$ for every test sample from the unseen client.

3.4 ERROR BOUND ANALYSIS ON UNSEEN CLIENTS WITH TSMOBN

In this section, we show that using TsmoBN yields a lower bound on generalization error compared to using training population statistics on unseen clients. Following the assumption that data points are sampled from Gaussian distribution, we denote the Gaussian distribution for client k as \mathcal{N}_k and the Gaussian distribution for the unseen client as \mathcal{N}_u . Based on the generalization bound for unseen domains given by Albuquerque et al. (2019), we show the error bound for a network with BN layers.

Theorem 3.2 (generalization bound extended from Albuquerque et al. (2019) for unseen client)

Let \mathcal{H} be a hypothesis class of a neural network with VC-dimension d and a sample of size n_k be drawn from Gaussian distribution for each seen clients and unseen clients in a FL system. Then, $\forall \alpha \in \mathbb{R}_+^n, \sum_{k=1}^K \alpha_k = 1$, with probability at least $1 - \delta$ over the choice of samples, for each $h \in \mathcal{H}$,

$$\epsilon_u(h_{\mathcal{N}_u}) \leq \epsilon_{\tilde{\mathcal{N}}} \left(\sum_{k=1}^K \alpha_k h_{\mathcal{N}_k} \right) + \sum_{k=1}^K \alpha_k \left(\frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{N}_k, \mathcal{N}_u) + \lambda_k \right) + 4 \sqrt{\frac{2d \log(2n_k K) + \log(4/\delta)}{n_k K}},$$

where $\epsilon_{\mathcal{N}}(h) = \mathbb{E}_{\mathbf{x} \sim \mathcal{N}}[|h(\mathbf{x}) - y|]$, h is a hypothesis function, $\tilde{\mathcal{N}}$ is the mixture of training samples with size $K n_k$, $d_{\mathcal{H}\Delta\mathcal{H}}$ measures distance between two distributions, $\lambda_k = \epsilon_{\mathcal{N}_i}(h^*) + \epsilon_{\mathcal{N}_u}(h^*)$, and $h^* = \arg \min_{h \in \mathcal{H}} \epsilon_{\mathcal{N}_i}(h) + \epsilon_{\mathcal{N}_u}(h)$.

The proof and details about the definition of the notations are given in Appendix C.1. This Theorem 3.2 gives the error bound of an unseen client for the vanilla FL. Based on this error bound, our proposed TsmoBN using surrogate variable further decreases the distances between distributions of the seen and unseen clients, i.e. after the intervention, $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{N}_i, \tilde{\mathcal{N}}_u) < d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{N}_i, \mathcal{N}_u)$, where $\tilde{\mathcal{N}}_u$ denotes the normalized feature distribution for the unseen client after intervention. Considering the $(\mathbb{E}[\mu_u], \mathbb{E}[\sigma_u^2])$ as the approximated estimation of true mean and variance for the unseen client, using $(\mathbb{E}[\mu_u], \mathbb{E}[\sigma_u^2])$ to normalize raw features R of unseen client reduces the divergence of the distributions of the normalized features between seen and unseen clients. More specifically, when the assumption that input data are all Gaussian distributed holds and assuming that \mathcal{H} is a class of two-layer neural network with batch normalization, the divergence becomes zero. We give a formal statement and proofs for zero-divergence error bound in Appendix C.2.

4 EXPERIMENTS

We validate our proposed solution TsmoBN for the unseen client generalization problem on three datasets. We conduct extensive experiments to show that: 1) our method can improve the generalizability of state-of-the-art FL algorithms on unseen clients, 2) how the results would be affected with ablation studies on four different key points, 3) our approach is orthogonal to current domain generalization methods and can be readily incorporated to gain a further performance boost.

4.1 DATASETS AND EXPERIMENTAL SETTINGS

Digit classification: We use the public Digits-DG (Zhou et al., 2020) dataset which includes images of ten digits, from 4 domains with drastic differences in font style, color, and background. Each domain is regarded as a client, and we use leave-one-client-out validation (i.e., three clients for training and one client for testing), with each client treated once as in D^u .

Methods	noise			blur				weather			digital				
	gaussian	shot	impulse	defocus	glass	motion	zoom	snow	frost	fog	bright	contrast	elastic	pixelate	jpeg
FedAvg	58.30	71.09	30.46	96.17	66.46	92.58	88.50	95.07	95.16	96.34	96.30	95.76	91.31	95.92	94.13
+TsmoBN +9.62	92.99	94.20	79.68	96.29	87.33	95.79	95.70	95.77	96.02	96.38	96.24	96.20	94.37	96.11	94.74
FedProx	58.42	71.21	31.68	96.22	64.67	91.83	87.28	95.00	95.04	96.34	96.35	95.70	90.69	95.98	93.88
+TsmoBN +9.72	92.65	94.15	79.30	96.34	86.37	95.70	95.65	95.75	95.99	96.42	96.28	96.26	94.36	96.16	94.74
FedNova	75.82	84.79	32.62	96.69	75.69	94.85	92.88	95.78	95.47	96.66	96.83	96.42	94.43	96.59	95.75
+TsmoBN +7.21	95.55	96.09	86.10	96.75	91.37	96.24	96.20	96.39	96.49	96.76	96.81	96.71	95.53	96.64	95.82
FedAdam	58.25	70.78	32.41	96.37	67.97	92.81	88.85	93.70	93.62	95.63	96.41	95.87	90.87	96.05	93.92
+TsmoBN +9.71	93.24	94.53	80.40	96.48	87.58	95.89	95.85	95.11	95.60	96.15	96.45	96.37	94.38	96.23	94.87
FedBN	55.87	68.84	28.44	96.04	64.56	90.38	84.95	94.38	94.42	96.06	96.11	94.61	87.26	95.47	91.99
+TsmoBN +10.60	91.53	93.01	77.30	96.42	85.48	95.71	95.52	95.46	95.87	96.35	96.30	96.17	93.48	96.06	93.77

Table 1: Performance of TsmoBN for generalization on CIFAR-10-C dataset. Each column represents an unseen client under leave-one-CorruptionType-out validation. *Results with the standard deviation are shown in Appendix E.2.2.*

CIFAR-10-C image recognition: This dataset includes CIFAR-10 images with 15 different styles of corruptions which are grouped into four types (i.e., noise, blur, weather, digital) (Hendrycks & Dietterich, 2019). We adopt its level-1 corruption severity as it is the closest to real-world situations. Each corruption is regarded as a client, and we use leave-one-CorruptionType-out validation (i.e., all corruption styles in three types form D^s , all corruption styles of the remaining type is D^u). This is designed to keep a certain corruption type be strictly unseen in training. The training client number is between 10 to 12. We use AlexNet (Krizhevsky et al., 2012) with BN added after each convolution and fully-connected layer (as did in Hendrycks & Dietterich (2019)).

Medical image diagnosis: We use the public tumor classification dataset Camelyon17, which shows histology images with different stains from 5 hospitals (Bandi et al., 2018). Following the setting of source/target domains in literature (Bandi et al., 2018; Koh et al., 2020), we use 3 hospitals for FL training, and test on 2 hospitals as unseen clients (denoted *hospital4* and *hospital5*). We use DenseNet121 (Huang et al., 2017). This dataset validates our method on a large-scale (i.e., 4.5 million images) real-world healthcare scenario, demonstrating its potential for practical applications.

All our FL models are trained with cross-entropy loss and SGD optimizer. As the sample sizes of sub-domains are not identical within each dataset, to avoid distracting the focus on the feature shift due to sample imbalance, we truncate the sample size of each client to their respective smallest number. We also demonstrate the effectiveness of our approach under imbalanced data setting in Appendix E.2.1. The data in D^s are randomly split to 80% training and 20% validation at each client. All data in the unseen client(s) are tested as unseen distributions. Without particular notice, we report the image classification accuracy results obtained with a batch size of 32 and a momentum value of 0.9. For complete experiment results and more implementation details (e.g., model architecture, image classification categories, pre-processings, learning rate, training details, computing resources, etc.), please refer to Appendix E.

4.2 IMPROVE UNSEEN CLIENT GENERALIZATION FOR STATE-OF-THE-ART FL ALGORITHMS

We implement our proposed TsmoBN on current state-of-the-art FL algorithms to show their performance improvement for unseen client generalization, including **FedAvg** (McMahan et al., 2017), **FedProx** (Li et al., 2020), **FedNova** (Wang et al., 2020a), **FedAdam** (Reddi et al., 2021) and **FedBN** (Li et al., 2021). For each FL algorithm, we compare two methods: 1) conventional inference on unseen clients with training-estimated BN statistics, 2) our proposed using test-specific momentum BN as surrogate variable intervention. Table 1, Table 2 and Fig. 2 respectively show the results on three datasets, which demonstrate that our proposed TsmoBN can consistently boost unseen client generalization performance for all these FL algorithms on all datasets, validating its general efficacy. Specifically, as shown in Table 2, our method improves over 3% accuracy on *svhn* compared to baselines, since *svhn* has very different appearances from others. All FL algorithms achieve similar results on *mnist* since images in *mnist* are clean and simple. However, our approach can still make further improvements. It is also noticed that, our approach achieves a significant increase of unseen client testing accuracy on the challenging CIFAR-10-C dataset which presents a wide scope of feature shift, and the real-world medical dataset where the image appearances are dramatically different

Methods	<i>svhn</i>	<i>mnistm</i>	<i>mnist</i>	<i>syn</i>	
FedAvg	73.47	69.90	98.71	89.51	
+TsmoBN	+1.86	77.56	71.10	99.13	90.68
FedProx	73.22	69.01	98.79	89.55	
+TsmoBN	+1.92	77.18	71.10	99.17	90.80
FedNova	76.36	73.84	98.83	91.49	
+TsmoBN	+1.13	79.63	74.56	99.18	91.64
FedAdam	73.96	68.89	98.68	89.09	
+TsmoBN	+1.60	77.02	70.72	99.04	90.23
FedBN	71.81	69.06	98.35	88.53	
+TsmoBN	+2.35	76.85	71.32	99.98	89.99

Table 2: Performance of TsmoBN for generalization on digit classification dataset. Each column represents an unseen client. We implement our TsmoBN solution on five state-of-the-art FL algorithms.

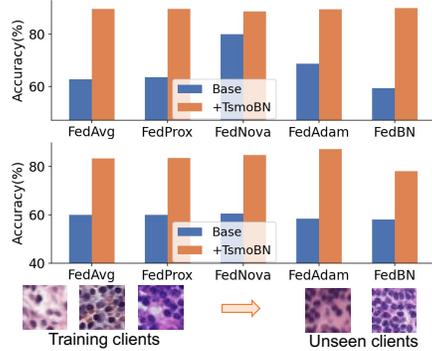


Figure 2: Results on medical image dataset with 3 training clients and 2 unseen clients. Blue bars are baseline FL results, orange bars are results of our TsmoBN.

across clients. On average for all unseen testing clients, our obtained unseen client generalization results increase at least 7.21% on CIFAR-10-C and 16.51% on the Camelyon17 medical dataset.

4.3 ABLATIONS STUDIES OF OUR METHOD

In particular, we conduct ablation studies using the medical dataset on four points including batch size, momentum value, fluctuation of statistics with different sampling order, and affine transformation in BN layers, which are important to implement our proposed method.

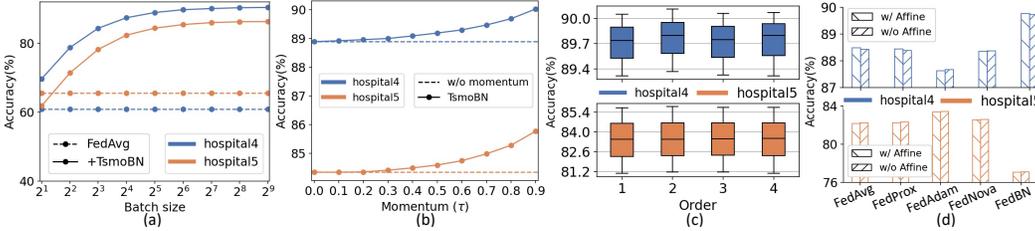


Figure 3: (a) and (b) study on how the batch size and momentum values would affect TsmoBN. (c) is the performance comparison of TsmoBN using different sampling sequences. (d) analyzes the effect of affine transformation parameters during testing.

One-batch experiment with different batch sizes: We conduct the one-batch study which explores how batch size can affect FL model performance on unseen clients with only one batch of samples (i.e., no mean or variance value tracked with momentum in TsmoBN). As shown in Fig. 3(a), we use FedAvg as the backbone and test the model on entire data of unseen clients with different batch sizes (from 2 to 512). The baseline (dotted line) using fixed training-estimated BN is not influenced by the batch size, but for our proposed TsmoBN, as expected, the test accuracy increases as we enlarge the batch size, because more samples benefit representative estimation of the statistical distributions.

Effects of momentum value: We further study whether the FL model performance is sensitive to the momentum value of τ , as results are shown in Fig. 3(b). In this experiment, we set the test batch size as 32, and change the value of momentum between $[0.0, 0.9]$ (noting that $\tau = 0$ means no historical values are tracked). It is observed that, within a wide range, using momentum is beneficial to performance improvement, especially when $\tau > 0.5$, while the overall performance is not very sensitive to the specific value of momentum (with accuracy changing within around 1%).

Sensitivity of sampling order: To study if the sampling order would affect the estimation during test time, we use the box plot to demonstrate the test accuracy with different sampling orders. For each order, we report the results of models trained with different random seeds. From Fig. 3(c) can be observed that our proposed TsmoBN has only slight changes under different orders, showing the potential under real applications that have no specific sampling sequence.

Affine transform in batch normalization during test: Affine transform in BN layers is an extra operation to enhance the representation capability and is continually updated during training. Here, we demonstrate that the key factor of TsmoBN is the test-specific statistics (i.e., the surrogate variable S in the test SCM) and our method is not sensitive to the affine transform. We use models trained with affine transforms in BN layers and test them with and without affine transformation respectively, as shown in Fig. 3(d). It is observed that, under different FL algorithms, removing the affine transform parameters almost makes no degeneration on the final model performance.

4.4 ENHANCE DOMAIN GENERALIZATION METHODS AS AN ORTHOGONAL APPROACH

Furthermore, we present the potential of our method on making additive improvements on FL-adaptable domain generalization methods, including **JiGen** (Carlucci et al., 2019), **CSD** (Piratla et al., 2020) and **RandConv** (Xu et al., 2021). We implemented FedAvg version of these methods on all three datasets, with results shown in Table 3, Table 4 and Fig. 4. It can be observed that the three domain generalization methods can generally bring performance improvement for conventional FedAvg for unseen client generalization. On top of that, our TsmoBN is orthogonal to these methods and can further boost their performance on all datasets. This shows the effectiveness and flexibility of our testing method without affecting the FL training paradigm.

Methods	<i>svhn</i>	<i>mnistm</i>	<i>mnist</i>	<i>syn</i>
FedAvg	73.47	69.90	98.71	89.51
JiGen	69.92	68.64	98.63	89.53
+TsmoBN +1.30	72.83	69.99	98.79	90.33
CSD	74.01	71.15	98.96	90.84
+TsmoBN +1.92	78.42	73.07	99.32	91.81
RandConv	76.55	80.53	98.95	92.97
+TsmoBN +2.50	82.39	83.92	99.39	93.30

Table 3: Results of adding our proposed TsmoBN to current domain generalization methods on digital classification dataset.

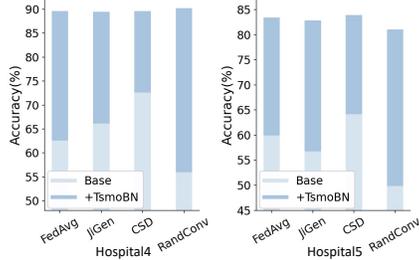


Figure 4: Results of adding our proposed TsmoBN to current domain generalization methods on medical dataset.

Methods	<i>noise</i>			<i>blur</i>				<i>weather</i>			<i>digital</i>				
	gaussian	shot	impulse	defocus	glass	motion	zoom	snow	frost	fog	bright	contrast	elastic	pixelate	jpeg
FedAvg	58.30	71.09	30.46	96.17	66.46	92.58	88.50	95.07	95.16	96.34	96.30	95.76	91.31	95.92	94.13
JiGen	61.13	70.51	33.86	95.24	66.29	88.88	84.97	94.44	94.84	96.40	95.29	92.63	85.11	93.91	89.21
+TsmoBN +7.01	87.23	88.79	74.91	94.58	77.84	92.62	92.45	94.12	94.94	95.97	93.66	93.35	87.82	92.36	87.21
CSD	64.71	77.05	32.99	96.38	71.35	93.64	90.45	95.48	95.48	96.48	96.40	95.85	91.60	96.11	94.42
+TsmoBN +8.46	93.97	94.92	82.01	96.52	88.36	96.06	95.85	96.08	96.20	96.44	96.43	96.43	94.70	96.23	95.05
RandConv	53.30	67.21	55.16	77.23	41.89	66.40	62.38	70.32	69.93	77.13	73.64	70.60	63.86	74.63	72.30
+TsmoBN +22.90	91.27	91.32	81.63	92.10	79.91	90.09	89.94	90.71	92.45	93.31	91.30	91.25	85.04	90.75	88.33

Table 4: Results of adding our proposed TsmoBN to current domain generalization methods on CIFAR-10-C dataset, which demonstrates that our solution is orthogonal to these methods.

5 CONCLUSION

This work tackles the unseen client generalization problem in FL by a novel causal view. We propose SCMs which explain the challenge of generalization problems in distributed learning paradigm, and further inspire the idea of using test-specific momentum batch normalization (TsmoBN) to tackle this generalization problem. We theoretically analyze TsmoBN within our test SCM from a view that test-specific statistics can be seen as a surrogate variable for causal intervention and show that TsmoBN yields improved generalization bounds by reducing the divergence of distributions between training and test. Our extensive experiments have demonstrated significant improvements for unseen client generalization for state-of-the-art FL algorithms, as well as orthogonal benefits for current domain generalization methods. Our study demonstrates that causal explanation can be successfully applied in federated learning, showing the great potential for future investigations beyond generalizability such as trustworthiness and safety.

6 ETHICS STATEMENT

Deploying models on data that differs from training samples is an ongoing challenge to ensure AI safety. Our work provides a promising solution to improve out-of-federation generalization for federated learning, which may be applied in various applications. For safety-critical tasks, however, such as disease diagnosis and autonomous vehicles, one needs to capture the unintended application-specific risk that cannot be mitigated by our interventional TsmoBN strategy.

7 REPRODUCIBILITY STATEMENT

For our proposed TsmoBN approach, a link to a anonymous downloadable source code can be found in Appendix A. The detailed algorithm box is shown in Appendix D. For theoretical results, assumptions and complete proofs are included in Sec. 3.3, Sec. 3.4, and Appendix B, C. For datasets used in the experiments, complete experiment settings, results, and implementation details, please refer to Sec. 4.1 and Appendix E.

REFERENCES

- Isabela Albuquerque, João Monteiro, Mohammad Darvishi, Tiago H Falk, and Ioannis Mitliagkas. Generalizing to unseen domains via distribution matching. *arXiv preprint arXiv:1911.00804*, 2019.
- Mathieu Andreux, Jean Ogier du Terrail, Constance Beguier, and Eric W Tramel. Siloed federated learning for multi-centric histopathology datasets. In *Domain Adaptation and Representation Transfer, and Distributed and Collaborative Learning*, pp. 129–139. Springer, 2020.
- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Ainesh Bakshi, Rajesh Jayaram, and David P Woodruff. Learning two layer rectified neural networks in polynomial time. In *Conference on Learning Theory (COLT)*, pp. 195–268, 2019.
- Peter Bandi, Oscar Geessink, Quirine Manson, Marcory Van Dijk, Maschenka Balkenhol, Meyke Hermesen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, et al. From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE Transactions on Medical Imaging*, 2018.
- Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2229–2238, 2019.
- Sitan Chen, Adam R. Klivans, and Raghu Meka. Learning deep relu networks is fixed-parameter tractable. *arXiv preprint arXiv:2009.13512*, 2020.
- Qi Dou, Daniel C. Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. In *Advances in Neural Information Processing Systems*, 2019.
- Ellery Eells. *Probabilistic causality*, volume 1. Cambridge University Press, 1991.
- Rong Ge, Jason D. Lee, and Tengyu Ma. Learning one-hidden-layer neural networks with landscape design. In *International Conference on Learning Representations*, 2018.
- Sander Greenland, Judea Pearl, and James M Robins. Causal diagrams for epidemiologic research. *Epidemiology*, pp. 37–48, 1999.
- Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*, 2020.
- Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019.
- Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- Xin Jin, Cuiling Lan, Wenjun Zeng, Zhibo Chen, and Li Zhang. Style normalization and restitution for generalizable person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3143–3152, 2020.
- Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J Reddi, Sebastian U Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for on-device federated learning. *arXiv preprint arXiv:1910.06378*, 2019.
- Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Bal-subramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. *arXiv preprint arXiv:2012.07421*, 2020.

- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018a.
- Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5400–5409, 2018b.
- Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. In *Conference on Machine Learning and Systems*, 2020.
- Xiaoxiao Li, Meirui Jiang, Xiaofei Zhang, Michael Kamp, and Qi Dou. FedBN: Federated learning on non-IID features via local batch normalization. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=6YEQUn0QICG>.
- Yuanzhi Li and Yang Yuan. Convergence analysis of two-layer neural networks with ReLU activation. In *Advances in neural information processing systems*, pp. 597–607, 2017.
- Quande Liu, Qi Dou, Lequan Yu, and Pheng Ann Heng. Ms-net: Multi-site network for improving prostate segmentation with heterogeneous mri data. *IEEE Transactions on Medical Imaging*, 2020.
- Divyat Mahajan, Shruti Tople, and Amit Sharma. Domain generalization using causal matching. In *International Conference on Machine Learning*, pp. 7313–7324. PMLR, 2021.
- Margaret Mooney Marini and Burton Singer. Causality in the social sciences. *Sociological methodology*, 18:347–409, 1988.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pp. 1273–1282, 2017.
- Nick Pawlowski, Daniel C Castro, and Ben Glocker. Deep structural causal models for tractable counterfactual inference. *arXiv preprint arXiv:2006.06485*, 2020.
- Judea Pearl. *Causality*. Cambridge university press, 2009.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.
- Krishna Pillutla, Sham M Kakade, and Zaid Harchaoui. Robust aggregation for federated learning. *arXiv preprint arXiv:1912.13445*, 2019.
- Vihari Piratla, Praneeth Netrapalli, and Sunita Sarawagi. Efficient domain generalization via common-specific low-rank decomposition. In *International Conference on Machine Learning*, pp. 7728–7738, 2020.
- Sashank J. Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and Hugh Brendan McMahan. Adaptive federated optimization. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=LkFG31B13U5>.
- Jacob C Reinhold, Aaron Carass, and Jerry L Prince. A structural causal model for mr images of multiple sclerosis. *arXiv preprint arXiv:2103.03158*, 2021.
- Amirhossein Reiszadeh, Farzan Farnia, Ramtin Pedarsani, and Ali Jadbabaie. Robust federated learning: The case of affine distribution shifts. *arXiv preprint arXiv:2006.08907*, 2020.
- Kaihua Tang, Jianqiang Huang, and Hanwang Zhang. Long-tailed classification by keeping the good and removing the bad momentum causal effect. *Advances in Neural Information Processing Systems*, 33, 2020.

- Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in Neural Information Processing Systems*, 33, 2020a.
- Rongguang Wang, Pratik Chaudhari, and Christos Davatzikos. Harmonization with flow-based causal inference. *arXiv preprint arXiv:2106.06845*, 2021a.
- Tan Wang, Chang Zhou, Qianru Sun, and Hanwang Zhang. Causal attention for unbiased visual recognition. *arXiv preprint arXiv:2108.08782*, 2021b.
- Wei Wang, Haojie Li, Zhengming Ding, and Zhihui Wang. Rethink maximum mean discrepancy for domain adaptation. *arXiv preprint arXiv:2007.00689*, 2020b.
- Herman Wold. Causality and econometrics. *Econometrica: Journal of the Econometric Society*, pp. 162–177, 1954.
- Zhenlin Xu, Deyi Liu, Junlin Yang, Colin Raffel, and Marc Niethammer. Robust and generalizable visual representation learning via random convolutions. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=BVSM0x3EDK6>.
- Tianyi Zhang, Ikko Yamane, Nan Lu, and Masashi Sugiyama. A one-step approach to covariate shift adaptation. In *Asian Conference on Machine Learning*, pp. 65–80. PMLR, 2020.
- Shanshan Zhao, Mingming Gong, Tongliang Liu, Huan Fu, and Dacheng Tao. Domain generalization via entropy regularization. *Advances in Neural Information Processing Systems*, 33, 2020.
- Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.
- Kai Zhong, Zhao Song, Prateek Jain, Peter L. Bartlett, and Inderjit S. Dhillon. Recovery guarantees for one-hidden-layer neural networks. In *International Conference on Machine Learning*, 2017.
- Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Learning to generate novel domains for domain generalization. In *European Conference on Computer Vision*, pp. 561–578. Springer, 2020.

A NOTATION TABLE

Notations	Description
\mathcal{X}	input space $\mathcal{X} \subseteq \mathbb{R}^d$
\mathcal{Y}	label space $\mathcal{Y} \in \{0, 1\}$
X	inputs variable, which denotes samples $\mathbf{x} \in \mathcal{X}$
Y	label variable of the inputs variable X
\mathcal{D}_k^s	the input data distribution for the k -th training client
\mathcal{D}^u	data distribution for an unseen client
K	number of clients
n_k	size of samples drawn from distribution for the k -th client
$p(\mathbf{x}, y)$	density of sample pairs from a specific distribution
R	raw extracted feature, i.e., the output of middle layers in neural network
F	normalized feature (i.e., the outputs of batch normalization layers)
\mathcal{G}^{tr}	directed acyclic graph representation of structural causal model for training
\mathcal{G}^{te}	directed acyclic graph representation of structural causal model for testing
S	surrogate variable, i.e., mean and variance calculated from R during testing
$P(\cdot)$	probability distribution
\mathcal{V}	variables in a certain structural causal model
PA_X	parent vertices of input X
U	unobserved variables
$\perp\!\!\!\perp$	notation for independent
μ	mean of data batch, a parameter of batch normalization layer
σ^2	variance of data batch, a parameter of batch normalization layer
τ	momentum value for causal intervention
\mathcal{H}	hypothesis class of two layer neural network with VC -dimension d
\mathcal{N}_i	gaussian distribution for the i th client
\mathcal{N}_u	gaussian distribution for the unseen client
$\tilde{\mathcal{N}}$	mixture of training samples with size Kn_k
λ_i	error of optimal hypothesis h^* on seen and unseen clients
$\tilde{\mathcal{N}}_u$	gaussian distribution for the unseen client after causal intervention
$\epsilon_u(h)$	risk term for unseen client with hypothesis function h
$\epsilon_s(h)$	risk term for seen clients with hypothesis function h
$\epsilon_u(h_{\mathcal{N}_u})$	risk term for unseen client with hypothesis function $h_{\mathcal{N}_u}$ with gaussian distribution assumption
$\epsilon_{\tilde{\mathcal{N}}}(h_{\mathcal{N}_i})$	risk term for seen training clients with hypothesis function $h_{\mathcal{N}_i}$ with gaussian distribution assumption
$d_{\mathcal{H}}(\cdot, \cdot)$	distance between two distributions

Table A.1: Notations occurred in the paper.

Code and data availability: We have included the example code to demonstrate our proposed method, please find in Google Drive via the below anonymous link <https://drive.google.com/file/d/1t9JOhegUZRHIhMkGaX8b0VYd5zpsPTFj/view?usp=sharing>

B CAUSAL ANALYSIS AND PROOFS

B.1 INFEASIBILITY OF BACK-DOOR ADJUSTMENT IN DG

In this section, we first show the formal statement on the infeasibility of backdoor adjustment for domain generalization in FL, and then give the proofs to get our proposed proposition.

Proposition B.1 (Infeasibility of back-door adjustment) *Consider an SCM over variables \mathcal{V} containing input X , parents PA_X of X , label Y and unobserved variables U in federated learning, the causal effect $P(Y|do(X))$ is identifiable whenever $\{X \cup Y \cup PA_X\} \subseteq V \setminus U$, i.e. whenever X , Y , and all parents of variables in X are measured. The back-door adjustment $P(Y|do(X)) = \sum_{PA_X} P(Y|X, PA_X)P(PA_X)$ fails whenever PA_X is not allowed to be accessed.*

To conduct intervention on a certain vertex X in structural causal model (SCM) via back-door adjustment, the key idea is summing the conditional probability over a set of parent vertices of X . Based on Chapter 6 Peters et al. (2017), we first define the valid adjustment set as following:

Definition B.2 (Valid adjustment set) *Consider an SCM over variables \mathcal{V} and let $Y \notin PA_X$ (otherwise we have $P(Y|do(X)) = P(Y)$). We call a set $\mathcal{D} \subseteq \mathcal{V} \setminus \{X, Y\}$ a valid adjustment set for the ordered pair (X, Y) if*

$$P(Y|do(X)) = \sum_{\mathbf{d}} P(Y|X, \mathbf{d})P(\mathbf{d}). \quad (6)$$

Here, the sum (could also be an integral) is over the range of \mathcal{D} , that is, over all values \mathbf{d} that \mathcal{D} can take.

From the valid adjustment set definition, we can see that the intervention $P(Y|do(X))$ can be achieved by adjusting over all values \mathbf{d} that \mathcal{D} can take. Before conducting the adjustment as shown in the right hand side of Eq.(6), we should make sure the left hand side intervention is identifiable, i.e., all variables in the SCM should be measured, otherwise, the causal intervention $P(Y|do(X))$ is not computable. We give a theorem based on Chapter 3 Pearl (2009) to illustrate the identifiability and the relationships between identifiability and adjustment as following:

Theorem B.3 (Identifiability) *Consider an SCM over a subset \mathcal{V} of variables are measured, the causal effect $P(Y|do(X))$ is identifiable whenever $\{X \cup Y \cup PA_X\} \subseteq V$, that is, whenever X , Y , and all parents of variables in X are measured. The expression $P(Y|do(X))$ is then obtained by adjusting for PA_X as in Eq.(6).*

From the Definition B.2 and Theorem B.3, we can summarize that to conduct a valid back-door adjustment for a pair (X, Y) , we should satisfy two conditions:

1. The causal effect of intervention $P(Y|do(X))$ is *identifiable*
2. The set of parent variables PA_X for adjustment are *valid*.

In conventional domain generalization methods, the intervention $P(Y|do(X))$ is *identifiable*, since the intervention can be computed from the observational distribution and the graph structure, i.e., the graph representation of SCM is explicitly defined and all *variables* are measured. As there is no strict on data access, all data distributions in \mathcal{D} can be accessed. Hence, two conditions are satisfied, making the intervention $P(Y|do(X))$ certainly identifiable and can be computed via back-door adjustment. Considering our federated learning system, all variables are measured under the federated learning paradigm, so the causal effect $P(Y|do(X))$ is identifiable. However, when we want to achieve the intervention by adjusting parent variables (i.e., \mathcal{D}_S) of X , since each client data is strictly kept locally, we cannot find a valid adjustment set for (X, Y) . The second condition does not hold when we directly apply conventional domain generalization methods into federated learning, making the back-door adjustment failed. Summarizing the above analysis, we have the proposition B.1.

B.2 PROOF OF THEOREM 3.1

In this section we show how to achieve the intervention on F via intervening the surrogate variable S . The main proof sketch is using calculus of intervention to exchange the effect of intervening F and S .

First, we give a definition of d-separation based on Pearl (2009) to help understand the calculus of intervention.

Definition B.4 (d-separation) *Let A, B, C be the three non-intersecting subsets of variables in a directed acyclic graph \mathcal{G} . For any path between two variables, a collider is a variable where arrows of the path meet head-to-head. A path from A to B is said to be blocked by C if either a non-collider on the path is in C , or there is a collider on the path and neither the collider nor its descendants are in C . If all paths from A to B are blocked, then A is d-separated from B by C , denoted by $A \perp\!\!\!\perp B \mid C$.*

The d-separation helps to demonstrate the independency between variables in the SCM. If two variables are independent, there is no causal effects between them, indicating that the intervention on one variable will not have effects on the other one. Based on the d-separation, we can show different effects of intervention towards dependent or independent variables, so according to Peters et al. (2017), we introduce the rules of intervention calculus as following:

Theorem B.5 (Rules of intervention calculus) *Given a DAG \mathcal{G} and disjoint subsets of variables X, Y, Z and W , let $P(\cdot)$ stand for the probability distribution induced by that model, we have the following rules.*

Rule 1 (Insertion/deletion of observations):

$$P(Y \mid Z, W, do(X)) = P(Y \mid W, do(X)), \quad \text{if } (Y \perp\!\!\!\perp Z \mid X, W)_{\mathcal{G}_{\overline{X}}}$$

Rule 2 (Intervention/observation exchange):

$$P(Y \mid do(Z), W, do(X)) = P(Y \mid Z, W, do(X)), \quad \text{if } (Y \perp\!\!\!\perp Z \mid X, W)_{\mathcal{G}_{\overline{XZ}}}$$

Rule 3 (Insertion/deletion of interventions):

$$P(Y \mid do(Z), Wdo(X)) = P(Y \mid W, do(X)), \quad \text{if } (Y \perp\!\!\!\perp Z \mid X, W)_{\mathcal{G}_{\overline{X}, \overline{Z}(W)}}$$

Start from our formed SCM for testing, the test-specific batch-statistics variable S (i.e., mean and variance of batch data in testset) is introduced and can be served as surrogate variable with following conditions satisfied: (1) F intercepts all directed paths from S to Y ; and (2) $P(Y|do(S))$ is identifiable (according to Theorem B.3). The first condition makes sure the causal effect between S and Y is indirect, i.e., S affects Y totally through F , otherwise, the effects of intervening S and F will both be passed to Y directly, we can not eliminate either one intervention. The second condition ensures the intervention $P(Y|do(S))$ can be computed over our SCM.

Firstly, the *Rule 3* of Theorem B.5 can be applied on original intervention $P(Y|do(F))$, having

$$P(Y|do(F)) = P(Y|do(F), do(S)), \quad \text{because } (Y \perp\!\!\!\perp S \mid F)_{\mathcal{G}_{\overline{FS}}} \quad (7)$$

The Eq.(7) demonstrates that, due to the intervention on F , variable F is not be affected by its parent variable S , so the external intervention $do(S)$ is independent from the label Y . With the independency between S and Y , we can add S holding any values into the conditional probability, having $P(Y|do(F), do(S))$.

Then, consider relations in $\mathcal{G}_{\overline{FS}}$, intervention on F can be removed based on *Rule 2*, having:

$$P(Y|do(F), do(S)) = P(Y|do(S)), \quad \text{because } (Y \perp\!\!\!\perp S \mid F)_{\mathcal{G}_{\overline{FS}}} \quad (8)$$

This equation shows the conditions for an external intervention $do(S)$ to have same effect on Y as the passive observation $S = s$ (here s is the mean and variance of test-specific data), since in $\mathcal{G}_{\overline{FS}}$ all paths from S to Y are eliminated. Combining Eq.(7) and Eq.(8), we finally conduct the intervention on F with the help of surrogate variable S ,

$$P(Y|do(F)) = P(Y|do(S)). \quad (9)$$

Combing the above conditions for introducing the surrogate variable S and rules to exchange intervention, we have the theorem as below:

Theorem B.6 *With SCM defined for testing, if a variable S served as surrogate variable with following conditions satisfied: (1) F intercepts all directed paths from S to Y ; (2) $P(Y|do(S))$ is identifiable; (3) $(Y \perp\!\!\!\perp S|F)_{\mathcal{G}_{\overline{SF}}^{te}}$; and (4) $(Y \perp\!\!\!\perp F|S)_{\mathcal{G}_{\overline{SE}}^{te}}$, then we have*

$$P(Y|do(F)) = P(Y|do(S)). \quad (10)$$

C ERROR BOUND ANALYSIS AND PROOFS

In this section, we first give the proofs for the Theorem 3.2, and then we introduce a corollary showing the new error bound with reduced divergence term.

C.1 PROOF OF THEOREM 3.2

First, we introduce notation for our error bound analysis on out-of-federation clients.

Notation: We define the client as a tuple with the specific distribution samples from \mathcal{D} and a deterministic labeling function g . We consider two types of clients, *seen* clients $\langle \mathcal{D}^s, g^s \rangle$ and a *unseen* client $\langle \mathcal{D}^u, g^u \rangle$. Moreover, we define a *hypothesis* $h : \mathcal{X} \rightarrow \mathcal{Y}$, such that $h \in \mathcal{H}$, where \mathcal{H} is a set of candidate hypothesis, and finally define the risk of hypothesis h with respect to the labeling function g on distribution \mathcal{D}^s as $\epsilon_s(h, g) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}^s} [|h(\mathbf{x}) - g(\mathbf{x})|]$. We denote the risk of hypothesis h on \mathcal{D}^s as $\epsilon_s(h)$. Parallely, the risk of hypothesis h on \mathcal{D}^u are denoted as $\epsilon_u(h)$.

Distance between two distributions: Let \mathcal{H} be a hypothesis class for input space \mathcal{X} , and $\mathcal{A}_{\mathcal{H}}$ be the set of subsets of \mathcal{X} that are the support of some hypothesis in \mathcal{H} , i.e. for each $h \in \mathcal{H}$, $\{\mathbf{x} : \mathbf{x} \in \mathcal{X}, h(\mathbf{x}) = 1\} \in \mathcal{A}_{\mathcal{H}}$. Then the distance between two distributions \mathcal{D} and \mathcal{D}' is defined as : $d_{\mathcal{H}}(\mathcal{D}, \mathcal{D}') := 2 \sup_{A \in \mathcal{A}_{\mathcal{H}}} |P_{\mathcal{D}}(A) - P_{\mathcal{D}'}(A)|$.

The symmetric difference space $\mathcal{H} \Delta \mathcal{H}$ is defined as : $\mathcal{H} \Delta \mathcal{H} := \{h(\mathbf{x}) \oplus h'(\mathbf{x}) | h, h' \in \mathcal{H}\}$, where \oplus is the XOR operation. $\mathcal{A}_{\mathcal{H} \Delta \mathcal{H}}$ is defined as the set of all sets A such that $A = \{\mathbf{x} : \mathbf{x} \in \mathcal{X}, h(\mathbf{x}) \neq h'(\mathbf{x})\}$ for some $h, h' \in \mathcal{H}$ and it follows that a distance $d_{\mathcal{H} \Delta \mathcal{H}}$ is well-defined. The *ideal hypothesis* h^* minimizes combined source and target risk: $h^* = \arg \min_{h \in \mathcal{H}} \epsilon_s(h) + \epsilon_u(h)$ and the error of h^* is defined as $\lambda := \epsilon_s(h^*) + \epsilon_u(h^*)$.

Based on Albuquerque et al. (2019), the error bound of the unseen client given one seen client is bounded by:

$$\epsilon_u(h) \leq \epsilon_s(h) + \frac{1}{2} d_{\mathcal{H} \Delta \mathcal{H}}(\mathcal{D}^s, \mathcal{D}^u) + 4 \sqrt{\frac{2d \log(2n_k) + \log(4/\delta)}{n_k}} + \lambda. \quad (11)$$

Consider federated learning with K seen clients and an unseen client, samples size of each client is n_k , the upper bound of $d_{\mathcal{H} \Delta \mathcal{H}}$ can be further derived as below:

$$\begin{aligned} d_{\mathcal{H} \Delta \mathcal{H}}(\mathcal{D}^s, \mathcal{D}^u) &= 2 \sup_{A \in \mathcal{A}_{\mathcal{H} \Delta \mathcal{H}}} |P_{\mathcal{D}^s}(A) - P_{\mathcal{D}^u}(A)| \\ &= 2 \sup_{A \in \mathcal{A}_{\mathcal{H} \Delta \mathcal{H}}} \left| \sum_{k=1}^K \alpha_k (P_{\mathcal{D}_k^s}(A) - P_{\mathcal{D}^u}(A)) \right| \\ &\leq 2 \sup_{A \in \mathcal{A}_{\mathcal{H} \Delta \mathcal{H}}} \sum_{k=1}^K \alpha_k (|P_{\mathcal{D}_k^s}(A) - P_{\mathcal{D}^u}(A)|) \\ &\leq 2 \sum_{k=1}^K \alpha_k \sup_{A \in \mathcal{A}_{\mathcal{H} \Delta \mathcal{H}}} (|P_{\mathcal{D}_k^s}(A) - P_{\mathcal{D}^u}(A)|) \\ &= \sum_{k=1}^K \alpha_k d_{\mathcal{H} \Delta \mathcal{H}}(\mathcal{D}_k^s, \mathcal{D}^u), \end{aligned}$$

the first inequality is derived by the triangle inequality. With the same triangle inequality, we can get $\lambda \leq \sum_{k=1}^K \lambda_k$ similarly. Note that $\epsilon_s(h_u) = \epsilon_s(\sum_{k=1}^K \alpha_k h_{s,k})$ for $\forall h_u \in \mathcal{H}$. Replace the λ , $\epsilon_s(h_u)$ and $d_{\mathcal{H} \Delta \mathcal{H}}$ into Eq.(11), we have:

$$\epsilon_u(h) \leq \epsilon_s\left(\sum_{k=1}^K \alpha_k h_{s,k}\right) + \sum_{k=1}^K \alpha_k (d_{\mathcal{H} \Delta \mathcal{H}}(\mathcal{D}_k^s, \mathcal{D}^u) + \lambda_k) + 4 \sqrt{\frac{2d \log(2n_k K) + \log(4/\delta)}{n_k K}} \quad (12)$$

With the assumption that the input data points follow Gaussian distribution for all clients, we denote the Gaussian distribution for client k as \mathcal{N}_k , the Gaussian distribution for unseen client as \mathcal{N}_u and the

mixture of training samples with size $n_k K$ as $\tilde{\mathcal{N}}$, we can change Eq.(12) into following form:

$$\epsilon_u(h_{\mathcal{N}_u}) \leq \epsilon_{\tilde{\mathcal{N}}} \left(\sum_{k=1}^K \alpha_k h_{\mathcal{N}_k} \right) + \sum_{k=1}^K \alpha_k \left(\frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{N}_k, \mathcal{N}_u) + \lambda_k \right) + 4 \sqrt{\frac{2d \log(2n_k K) + \log(4/\delta)}{n_k K}}. \quad (13)$$

Extend to multiple unseen clients: Given n_u unseen clients, let \mathcal{N}_j be the distribution for the j -th unseen client, $j = 1, \dots, n_u$, and let $\tilde{\mathcal{N}}_u$ be the mixture of unseen clients. Then, the error bound is

$$\begin{aligned} \epsilon_u(h_{\mathcal{N}_u}) &\leq \sum_{j=1}^{n_u} \epsilon_{u_j}(h_{\mathcal{N}_u}) \\ &\leq n_u \epsilon_{\tilde{\mathcal{N}}} \left(\sum_{k=1}^K \alpha_k h_{\mathcal{N}_k} \right) + \sum_{j=1}^{n_u} \sum_{k=1}^K \alpha_k \left(\frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{N}_k, \mathcal{N}_{u_j}) + \lambda_{kj} \right) \\ &\quad + 4 \sqrt{\frac{2d \log(2n_k K) + \log(4/\delta)}{n_k K}}, \end{aligned}$$

where λ_{kj} is the risk of the optimal hypothesis on the mixture of seen client \mathcal{N}_k and unseen client \mathcal{N}_j .

C.2 COROLLARY OF REDUCED ERROR BOUND USING TMBN

We first give the formal statement of the new error bound with reduced divergence term using our proposed TMBN as the causal intervention.

Corollary C.1 (generalization bound with TMBN as causal intervention) *Under the assumptions of Theorem 3.2, and assume the network is a two-layer networks with BN layers, the error with causal intervention is bounded by*

$$\epsilon_u(h_{\tilde{\mathcal{N}}_u}) \leq \epsilon_{\tilde{\mathcal{N}}} \left(\sum_{k=1}^K \alpha_k h_{\mathcal{N}_k} \right) + \sum_{k=1}^K \alpha_k \lambda_k + 4 \sqrt{\frac{2d \log(2n_k K) + \log(4/\delta)}{n_k K}}.$$

The proof is given below, and the key idea is that, by using our approach TMBN as causal intervention, i.e., using the test statistics to normalize features, the divergence of the distributions of the normalized features between seen and unseen clients decreases. More specifically, if the input data are all Gaussian distributed, our causal intervention is tend to obtain the standard Gaussian distribution same as feature distributions during training, thus the divergence becomes zero.

As we assume that \mathcal{H} is a class of two layer neural network with batch normalization, then the input data is transformed as

$$\mathbf{z} := (V[W\mathbf{x}])^{-1/2}(W\mathbf{x} - \mathbb{E}[W\mathbf{x}]),$$

where $W \in \mathbb{R}^{d \times M}$ and M is the number of neurons. Since \mathbf{x} is Gaussian distributed, the transformed layer \mathbf{z} follows \mathcal{N} , where \mathcal{N} is the standard Gaussian distribution with mean $\mathbf{0}$ and covariance matrix as identity matrix. For $\forall h \in \mathcal{H}$, there exists a function q such that $q(\mathbf{z}) = h(\mathbf{x}; W)$ and for the true labeling function g , let $q_t(\mathbf{z}) = g(\mathbf{x}; W_t)$, where W_t is the true model parameter. Thus,

$$\epsilon_{\mathcal{N}_i}(h) = \mathbb{E}_{\mathbf{x} \sim \mathcal{N}_i} [|h(\mathbf{x}) - g(\mathbf{x})|] = \mathbb{E}_{\mathbf{z} \sim \mathcal{N}} [|q(\mathbf{z}) - q_t(\mathbf{z})|] = \epsilon_{\mathcal{N}}(q).$$

For the seen clients, the risk becomes $\epsilon_{\mathcal{N}}(q)$ instead of $\epsilon_{\mathcal{N}_i}(h)$, that means the covariate shift is eliminated by batch normalization.

For unseen client with vanilla FedAvg method, the input data is transformed as $\mathbf{z}_v := \Sigma_s^{-1/2}(W\mathbf{x} - \mu_s)$, where μ_s and Σ_s are the mean and variance matrix calculated using the seen clients data. Let \mathcal{N}_v be the distribution for \mathbf{z}_v , then the error bound becomes

$$\epsilon_{\mathcal{N}_v}(q_u) \leq \epsilon_{\mathcal{N}} \left(\sum_{k=1}^K \alpha_k q_k \right) + \sum_{k=1}^K \alpha_k \left(\frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{N}, \mathcal{N}_v) + \lambda_k \right) + 4 \sqrt{\frac{2d \log(2n_k K) + \log(4/\delta)}{n_k K}}.$$

For unseen client with proposed causal intervention, the input data is transformed as $\mathbf{z}_I := \Sigma_t^{-1/2}(W \mathbf{x} - \mu_t)$, where μ_t and Σ_t are the mean and variance matrix calculated using the seen clients data. Then we have $\mathbf{z}_I \sim \mathcal{N}$, and the error bound becomes

$$\begin{aligned} \epsilon_{\mathcal{N}}(q_u) &\leq \epsilon_{\mathcal{N}}\left(\sum_{k=1}^K \alpha_k q_k\right) + \sum_{k=1}^K \alpha_k \left(\frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{N}, \mathcal{N}) + \lambda_k\right) + 4\sqrt{\frac{2d \log(2n_k K) + \log(4/\delta)}{n_k K}} \\ &= \epsilon_{\mathcal{N}}\left(\sum_{k=1}^K \alpha_k q_k\right) + \sum_{k=1}^K \alpha_k \lambda_k + 4\sqrt{\frac{2d \log(2n_k K) + \log(4/\delta)}{n_k K}}. \end{aligned}$$

The causal intervention reduces the divergence of distributions of the seen and unseen clients.

D TSMOBN ALGORITHM

We describe the details algorithm of our proposed TsmoBN as following Algorithm 1.

Algorithm 1 Algorithm of TsmoBN

Notations: For a neural network with batch normalization layers, $TsmoBN(x; \mu, \sigma^2)$ denotes our proposed normalization method on samples x with mean μ and variance σ^2 . M denotes incoming batches of data, where each batch has m samples $\{x_1, \dots, x_m\}$. τ denotes the momentum value.

```

1: During testing:
2: for each batch  $t = 1, \dots, M$  do
3:   if  $t = 1$  then
4:     Initialize  $\mu = \frac{1}{m} \sum_{i=1}^m x_i, \sigma^2 = \frac{1}{m} \sum_{i=1}^m (x_i - \mu)^2$ 
5:   else
6:      $\mu = \tau\mu + (1 - \tau)\frac{1}{m} \sum_{i=1}^m x_i, \sigma^2 = \tau\sigma^2 + (1 - \tau)\frac{1}{m} \sum_{i=1}^m (x_i - \mu)^2$ 
7:   end if
8:    $TsmoBN(\{x_i\}_{i=1}^m; \mu, \sigma^2)$ 
9: end for

```

E EXPERIMENTS

In this part we show more experimental results and implementation details. Sec. E.1 gives more details for implementation, including dataset details, model architectures, training and testing details. Sec. E.2 shows more experiment results, more specifically, Sec. E.2.1 reports the results of TsmoBN using imbalanced data in camelyon17 dataset, Sec. E.2.2 presents all results of table and figure appeared in our manuscript with standard deviation for multiple runs, Sec. E.2.3 further studies hyper-parameters on four types (i.e., *noise*, *blur*, *weather* and *digital*) in CIFAR-10-C, Sec. E.2.4 demonstrates that our proposed TsmoBN is not sensitive to different sampling orders and in Sec. E.2.5, we conduct more experiments on four more corruptions types (*speckle noise*, *gaussian blur*, *spatter* and *saturate*) provided by CIFAR-10-C to further show the generalizability improvements and effects of involving more various distributions.

E.1 EXPERIMENTAL DETAILS

In this section we present more details for three datasets (Digits-DG, CIFAR-10-C and Camelyon17), model architectures and model training and testing settings.

Digit classification: For our 10 digits classification experiment on Digits-DG dataset, we use images from four different domains as shown in Fig. E.1 and reshape all images into $28 \times 28 \times 3$. We use a six-layer Convolutional Neural Network (CNN) and its details are listed in Table. E.2. We train the model for 100 epochs with the learning rate of 0.01 and select the best model with the highest validation accuracy.

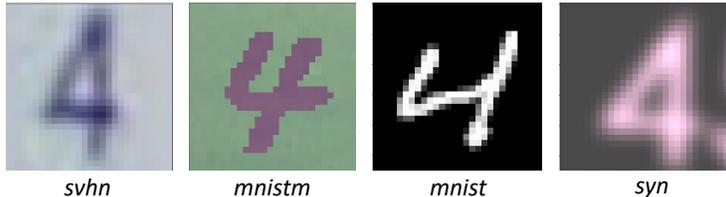


Figure E.1: Samples from four different domains in Digits-DG.

Layer	Details
1	Conv2D(3, 64, 5, 1, 2) BN(64), ReLU, MaxPool2D(2, 2)
2	Conv2D(64, 64, 5, 1, 2) BN(64), ReLU, MaxPool2D(2, 2)
3	Conv2D(64, 128, 5, 1, 2) BN(128), ReLU
4	FC(6272, 2048) BN(2048), ReLU
5	FC(2048, 512) BN(512), ReLU
6	FC(512, 10)

Table E.2: Model architecture of the Digits-DG experiment. For convolutional layer (Conv2D), we list parameters with sequence of input and output dimension, kernel size, stride and padding. For max pooling layer (MaxPool2D), we list kernel and stride. For fully connected layer (FC), we list input and output dimension. For BatchNormalization layer (BN), we list the channel dimension.

CIFAR-10-C image recognition: CIFAR-10-C contains ten categories of natural images with various corruption types, as shown in Fig. E.2 and we preprocess all images into $256 \times 256 \times 3$. We use adapted AlexNet Krizhevsky et al. (2012) pretrained on Imagenet dataset, and add BN layer after each convolutional layer and fully-connected layer (except the last layer), as the architecture shown in

Table E.3. We train the model for 200 epochs with learning rate of 0.001 and early stop the training if no improvements for validation accuracy over 10 epochs. The model is selected based on the highest validation accuracy.

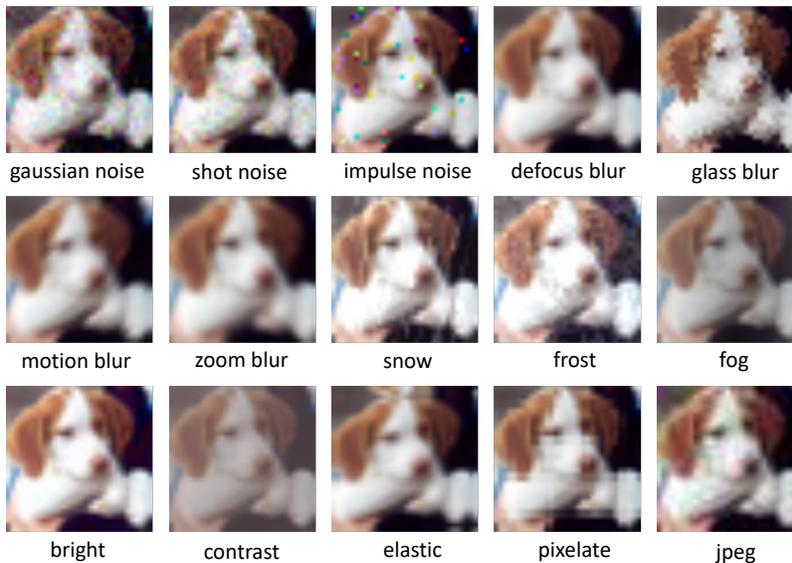


Figure E.2: Samples from fifteen different corruption types in CIFAR-10-C.

Layer	Details
1	Conv2D(3, 64, 11, 4, 2) BN(64), ReLU, MaxPool2D(3, 2)
2	Conv2D(64, 192, 5, 1, 2) BN(192), ReLU, MaxPool2D(3, 2)
3	Conv2D(64, 128, 5, 1, 2) BN(128), ReLU
4	Conv2D(384, 256, 3, 1, 1) BN(256), ReLU
5	Conv2D(256, 256, 3, 1, 1) BN(256), ReLU, MaxPool2D(3, 2)
6	AdaptiveAvgPool2D(6, 6)
7	FC(9216, 4096) BN(4096), ReLU
8	FC(4096, 4096) BN(4096), ReLU
9	FC(4096, 10)

Table E.3: Model architecture for CIFAR-10-C experiment. For convolutional layer (Conv2D), we list parameters with sequence of input and output dimension, kernel size, stride and padding. For max pooling layer (MaxPool2D), we list kernel and stride. For fully connected layer (FC), we list input and output dimension. For BatchNormalization layer (BN), we list the channel dimension.

Medical image diagnosis: For the histology image diagnosis, we use the all data from the public dataset Camelyon17 Bandi et al. (2018). Camelyon17 comprises 450,000 patches of breast cancer metastases in lymph node sections from 5 hospitals in the Netherlands. The task is to predict whether a given region of tissue contains any tumor tissue or not. All the data are preprocessed into shape of $96 \times 96 \times 3$. For neural network, we use the DenseNet121 Huang et al. (2017) and train it with learning rate of 0.001 for 200 epochs. We early stop the training if no improvements for validation accuracy over 10 epochs and select the model selected based on the highest validation accuracy.

For all experiments, we train and test our models on TITAN RTX GPU with Ubuntu18.04 operating system.

E.2 MORE EXPERIMENTS

Here we show more experimental results, including performance with imbalanced setting for real-world application (Sec. E.2.1), results reported in manuscript with standard deviation (Sec. E.2.2), and more hyper-parameter studies on batch size and momentum value (Sec. E.2.3), including four groups of corruption types (i.e., *noise*, *blur*, *weather* and *digital*) in CIFAR-10-C. In addition, in Sec. E.2.4 we show that our proposed TsmoBN is not sensitive to the order of testing batches. For CIFAR-10-C, we additionally show unseen client generalization results on extra 4 unseen clients (i.e., *speckle noise*, *gaussian blur*, *spatter* and *saturate*) in Sec. E.2.5.

E.2.1 IMBALANCED DATA SETTING FOR REAL-WORLD APPLICATION ON MEDICAL DATASET

In this section, we conduct experiments on camelyon17 dataset with imbalanced training data (i.e., original size of samples from each training client). From the Table E.4 can be observed that our proposed TsmoBN can effectively improve the performance with various existing FL algorithms even with the imbalanced training data across clients.

Methods	<i>hospital4</i>	<i>hospital5</i>
FedAvg	63.23(6.40)	63.31(8.10)
+TsmoBN +23.60	89.61(0.14)	84.13(0.16)
FedProx	62.45(6.71)	63.73(6.65)
+TsmoBN +24.29	89.86(0.41)	84.90(1.31)
FedNova	56.56(3.06)	58.16(9.80)
+TsmoBN +30.64	90.64(0.25)	85.35(2.19)
FedAdam	71.31(13.90)	62.03(6.34)
+TsmoBN +20.33	89.28(0.25)	84.71(0.24)
FedBN	68.06(16.01)	64.85(14.93)
+TsmoBN +17.41	89.39(1.20)	78.33(1.96)

Table E.4: Performance of TsmoBN under imbalanced data setting. We report results on medical image dataset with 3 training clients and 2 unseen clients. Each column represents an unseen client. We implement our approach on five state-of-the-art FL algorithms and report standard deviation (in bracket) for multiple runs.

E.2.2 DETAILED RESULTS WITH STANDARD DEVIATION FOR MULTIPLE RUNS

We run our experiments 3 times with random seed 0, 1, 2 and report the average results with standard deviation.

Improve unseen client generalization for state-of-the-art FL algorithms:

- For detailed version of digit classification results reported in Table 2, please refer to Table E.5.
- For detailed version of medical image diagnosis results reported in Table 1, please refer to Table E.6.
- For detailed version of CIFAR-10-C image recognition results reported in Fig. 2, please refer to Table E.7.

From results in three tables, we have the observation that besides generalization accuracy improvements, our proposed methods have lower standard deviation than baseline FL algorithms in general.

Methods	<i>svhn</i>	<i>mnistm</i>	<i>mnist</i>	<i>syn</i>
FedAvg	73.47(0.62)	69.90(0.82)	98.71(0.08)	89.51(0.14)
+TsmoBN +1.86	77.56(0.35)	71.10(0.24)	99.13(0.07)	90.68(0.14)
FedProx	73.22(0.48)	69.01(0.62)	98.79(0.10)	89.55(0.16)
+TsmoBN +1.92	77.18(0.35)	71.10(0.40)	99.17(0.10)	90.80(0.33)
FedNova	76.36(0.13)	73.84(0.57)	98.83(0.07)	91.49(0.49)
+TsmoBN +1.13	79.63(0.81)	74.56(0.49)	99.18(0.08)	91.64(0.41)
FedAdam	73.96(0.09)	68.89(0.63)	98.68(0.07)	89.09(0.20)
+TsmoBN +1.60	77.02(0.60)	70.72(0.52)	99.04(0.07)	90.23(0.05)
FedBN	71.81(0.51)	69.06(1.26)	98.35(0.18)	88.53(0.34)
+TsmoBN +2.35	76.85(0.25)	71.32(0.55)	99.98(0.12)	89.99(0.07)

Table E.5: Performance of TsmoBN for unseen client generalization on digit classification dataset. Each column represents an unseen client. We implement our approach on five state-of-the-art FL algorithms and report standard deviation (in bracket) for multiple runs.

Methods	<i>noise</i>			<i>blur</i>				<i>weather</i>			<i>digital</i>				
	gaussian	shot	impulse	defocus	glass	motion	zoom	snow	frost	fog	bright	contrast	elastic	pixelate	jpeg
FedAvg	58.30 (0.97)	71.09 (0.40)	30.46 (3.06)	96.17 (0.15)	66.46 (1.31)	92.58 (0.71)	88.50 (1.53)	95.07 (0.17)	95.16 (0.23)	96.34 (0.06)	96.30 (0.03)	95.76 (0.08)	91.31 (0.79)	95.92 (0.05)	94.13 (0.33)
+TsmoBN +9.62	92.99 (0.58)	94.20 (0.26)	79.68 (1.07)	96.29 (0.12)	87.33 (0.89)	95.79 (0.11)	95.70 (0.16)	95.77 (0.14)	96.02 (0.13)	96.38 (0.09)	96.24 (0.05)	96.20 (0.12)	94.37 (0.19)	96.11 (0.12)	94.74 (0.07)
FedProx	58.42 (3.85)	71.21 (2.94)	31.68 (2.83)	96.22 (0.12)	64.67 (1.95)	91.83 (1.14)	87.28 (2.10)	95.00 (0.17)	95.04 (0.12)	96.34 (0.14)	96.35 (0.05)	95.70 (0.16)	90.69 (0.59)	95.98 (0.06)	93.88 (0.19)
+TsmoBN +9.72	92.65 (0.69)	94.15 (0.63)	79.30 (1.17)	96.34 (0.23)	86.37 (0.85)	95.70 (0.03)	95.65 (0.11)	95.75 (0.12)	95.99 (0.12)	96.42 (0.11)	96.28 (0.15)	96.26 (0.06)	94.36 (0.14)	96.16 (0.09)	94.74 (0.14)
FedNova	75.82 (0.98)	84.79 (0.54)	32.62 (3.58)	96.69 (0.05)	75.69 (1.95)	94.85 (0.10)	92.88 (0.37)	95.78 (0.22)	95.47 (0.32)	96.66 (0.06)	96.83 (0.09)	96.42 (0.12)	94.43 (0.25)	96.59 (0.02)	95.75 (0.11)
+TsmoBN +7.21	95.55 (0.20)	96.09 (0.12)	86.10 (0.59)	96.75 (0.04)	91.37 (0.46)	96.24 (0.10)	96.20 (0.08)	96.39 (0.15)	96.49 (0.10)	96.76 (0.05)	96.81 (0.09)	96.71 (0.06)	95.53 (0.05)	96.64 (0.10)	95.82 (0.08)
FedAdam	58.25 (3.30)	70.78 (2.28)	32.41 (1.97)	96.37 (0.12)	67.97 (1.17)	92.81 (0.45)	88.85 (0.83)	93.70 (2.20)	93.62 (2.35)	95.63 (1.20)	96.41 (0.19)	95.87 (0.35)	90.87 (1.51)	96.05 (0.16)	93.92 (1.25)
+TsmoBN +9.71	93.24 (0.11)	94.53 (0.18)	80.40 (0.46)	96.48 (0.14)	87.58 (0.88)	95.89 (0.20)	95.85 (0.15)	95.11 (1.18)	95.60 (0.54)	96.15 (0.40)	96.45 (0.08)	96.37 (0.06)	94.38 (0.54)	96.23 (0.09)	94.87 (0.45)
FedBN	55.87 (4.83)	68.84 (4.55)	28.44 (1.59)	96.04 (0.19)	64.56 (1.59)	90.38 (1.13)	84.95 (1.41)	94.38 (0.25)	94.42 (0.43)	96.06 (0.14)	96.11 (0.22)	94.61 (1.05)	87.26 (2.51)	95.47 (0.57)	91.99 (1.48)
+TsmoBN +10.60	91.53 (1.17)	93.01 (1.02)	77.30 (1.08)	96.42 (0.13)	85.48 (0.63)	95.71 (0.30)	95.52 (0.21)	95.46 (0.14)	95.87 (0.07)	96.35 (0.11)	96.30 (0.13)	96.17 (0.24)	93.48 (1.00)	96.06 (0.28)	93.77 (0.87)

Table E.6: Performance of TsmoBN for unseen client generalization on CIFAR-10-C dataset. Each column represents an unseen client under leave-one-CorruptionType-out validation. We implement our approach on five state-of-the-art FL algorithms and report standard deviation (in bracket) for multiple runs.

Enhance domain generalization methods as an orthogonal approach:

- For detailed version of digit classification results reported in Table 3, please refer to Table E.8.
- For detailed version of medical image diagnosis results reported in Table 4, please refer to Table E.9.
- For detailed version of CIFAR-10-C image recognition results reported in Fig. 4, please refer to Table E.10.

By applying our proposed approach on domain generalization methods, especially for data with strong distribution changes (e.g., *noise* in CIFAR-10-C), we can see the standard deviation is mitigated.

Methods	<i>hospital4</i>	<i>hospital5</i>
FedAvg	62.62(5.82)	59.95(7.93)
+TsmoBN +25.27	89.66(0.40)	83.44(2.34)
FedProx	63.41(5.64)	59.99(8.75)
+TsmoBN +24.88	89.66(0.37)	83.51(2.58)
FedNova	80.02(9.41)	60.54(2.06)
+TsmoBN +16.51	88.72(1.13)	84.86(3.07)
FedAdam	68.71(10.46)	58.41(9.63)
+TsmoBN +24.77	89.46(0.37)	87.20(6.34)
FedBN	59.36(15.64)	58.14(2.99)
+TsmoBN +25.27	89.98(1.33)	78.06(2.65)

Table E.7: Performance of TsmoBN for unseen client generalization on medical image dataset with 3 training clients and 2 unseen clients. Each column represents an OOF client. We implement our approach on five state-of-the-art FL algorithms and report standard deviation (in bracket) for multiple runs.

Methods	<i>svhn</i>	<i>mnistm</i>	<i>mnist</i>	<i>syn</i>
FedAvg	73.47(0.62)	69.90(0.82)	98.71(0.08)	89.51(0.14)
JiGen	69.92(1.28)	68.64(0.12)	98.63(0.15)	89.53(0.23)
+TsmoBN +1.30	72.83(0.56)	69.99(0.16)	98.79(0.17)	90.33(0.26)
CSD	74.01(0.54)	71.15(0.59)	98.96(0.11)	90.84(0.28)
+TsmoBN +1.92	78.42(0.77)	73.07(0.39)	99.32(0.06)	91.81(0.11)
RandConv	76.55(1.19)	80.53(0.59)	98.95(0.09)	92.97(0.54)
+TsmoBN +2.50	82.39(0.41)	83.92(0.52)	99.39(0.02)	93.30(0.16)

Table E.8: Results of adding our proposed TsmoBN to current domain generalization methods on digital classification dataset. We report standard deviation (in bracket) for multiple runs.

Methods	<i>noise</i>			<i>blur</i>				<i>weather</i>			<i>digital</i>				
	gaussian	shot	impulse	defocus	glass	motion	zoom	snow	frost	fog	bright	contrast	elastic	pixelate	jpeg
JiGen	61.13 (0.96)	70.51 (0.92)	33.86 (1.02)	95.24 (2.08)	66.29 (3.93)	88.88 (3.74)	84.97 (4.38)	94.44 (0.23)	94.84 (0.08)	96.40 (0.02)	95.29 (0.64)	92.63 (1.35)	85.11 (1.00)	93.91 (0.87)	89.21 (1.48)
+TsmoBN +7.01	87.23 (0.81)	88.79 (0.74)	74.91 (0.83)	94.58 (2.58)	77.84 (3.01)	92.62 (3.18)	92.45 (3.51)	94.12 (0.10)	94.94 (0.19)	95.97 (0.07)	93.66 (0.98)	93.35 (1.16)	87.82 (1.24)	92.36 (1.09)	87.21 (1.58)
CSD	64.71 (1.64)	77.05 (0.28)	32.99 (1.24)	96.38 (0.07)	71.35 (1.18)	93.64 (0.50)	90.45 (1.25)	95.48 (0.10)	95.48 (0.13)	96.48 (0.18)	96.40 (0.06)	95.85 (0.18)	91.60 (0.91)	96.11 (0.14)	94.42 (0.37)
+TsmoBN +8.46	93.97 (0.34)	94.92 (0.21)	82.01 (0.67)	96.52 (0.09)	88.36 (0.66)	96.06 (0.06)	95.85 (0.04)	96.08 (0.03)	96.20 (0.06)	96.44 (0.06)	96.43 (0.08)	96.43 (0.15)	94.70 (0.11)	96.23 (0.11)	95.05 (0.18)
RandConv	53.30 (3.33)	67.21 (3.61)	55.16 (3.96)	77.23 (0.83)	41.89 (4.65)	66.40 (1.41)	62.38 (1.08)	70.32 (3.16)	69.93 (3.22)	77.13 (3.16)	73.64 (2.63)	70.60 (3.27)	63.86 (2.22)	74.63 (2.23)	72.30 (2.59)
+TsmoBN +22.90	91.27 (1.23)	91.32 (1.38)	81.63 (1.15)	92.10 (0.97)	79.91 (1.53)	90.09 (1.12)	89.94 (1.18)	90.71 (1.40)	92.45 (1.22)	93.31 (1.10)	91.30 (0.21)	91.25 (0.14)	85.04 (0.14)	90.75 (0.12)	88.33 (0.31)

Table E.9: Results of adding our proposed TsmoBN to current domain generalization methods on CIFAR-10-C dataset, which demonstrates that our solution is orthogonal to these methods. We report standard deviation (in bracket) for multiple runs.

E.2.3 HYPER-PARAMETER STUDIES ON FOUR CORRUPTION TYPES IN CIFAR-10-C

In addition to the results on medical dataset which are shown in Sec. 4.3, we further conduct studies on the four groups of corruptions (i.e., *blur*, *weather* and *digital*) and plot the results in Fig. E.3 and Fig. E.4.

One-batch experiment with different batch sizes: We first show accuracy changes with different batch size. The observations are consistent that the unseen client test accuracy increases as we enlarge

Methods	<i>hospital4</i>	<i>hospital5</i>
JiGen	66.10(4.95)	56.75(9.97)
+TsmoBN +24.75	89.50(0.41)	82.83(1.62)
CSD	72.60(9.17)	64.19(5.63)
+TsmoBN +18.37	89.60(0.14)	83.92(0.34)
RandConv	55.91(7.49)	49.85(0.32)
+TsmoBN +32.78	90.23(0.94)	81.10(5.25)

Table E.10: Results of adding our proposed TsmoBN to current domain generalization methods on medical dataset. We report standard deviation (in bracket) for multiple runs.

the batch size, because more samples benefit representative estimation of the statistical distributions. Meanwhile, we also see that the accuracy becomes saturated near to the highest performance after batch size is sufficiently large. For corruption types bringing strong distribution changes (e.g., *noise*, *blur*), our method outperforms baseline method (FedAvg) at very small batch size (e.g., batch size of 8). For corruption types with less changes as *weather* and *digital*, our methods achieve comparable or better results after batch size of 16, showing less requirements on large batch size.

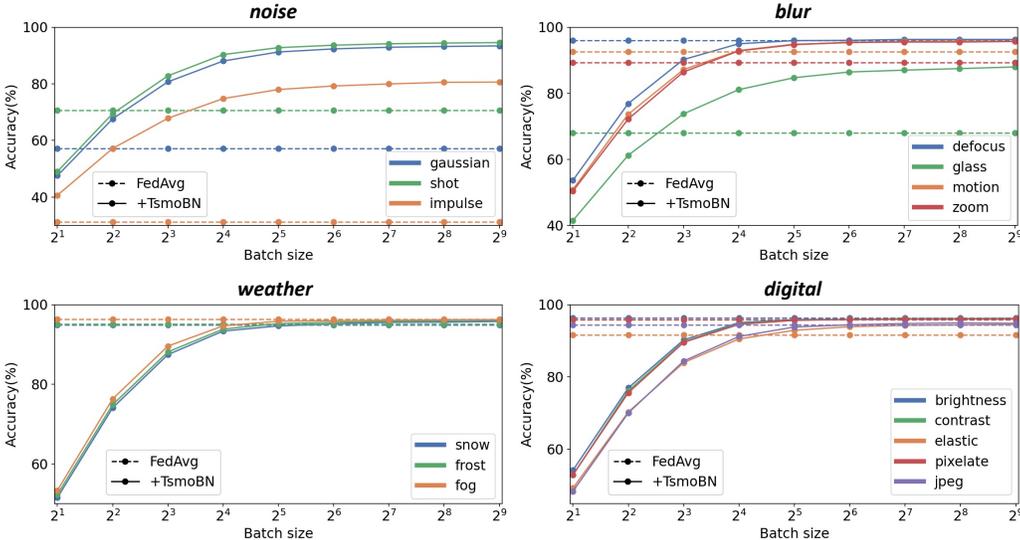


Figure E.3: Studies on how batch size would affect TsmoBN performance. We report FedAvg on three unseen corruption types (i.e., noise, blur, weather and digital) in CIFAR-10-C.

Studies on momentum value: For the momentum values studies, we set the testing batch size as 32, and change the value of momentum between [0.0, 0.9]. Please note that when momentum value is 0, it is equal to not using momentum. From the Fig. E.4 we can observe that, within a wide range, using momentum is beneficial to improve performance, while overall performance is not very sensitive to the specific value of momentum (with accuracy changing within around 1%).

E.2.4 SAMPLING ORDER SENSITIVITY STUDY ON TSMOBN

Besides results on the medical dataset camelyon17, we further validate our proposed TsmoBN on batches with different incoming orders to show the stability. Specifically, we shuffle the test set with 5 different random seeds (i.e., 0, 1, 2, 3, 4) and draw the accuracies as shown in Fig. E.5. We show results on four different types of corruptions (i.e., *noise*, *blur*, *weather* and *digital*) in CIFAR-10-C. The observation clearly shows the stability of our proposed TsmoBN, all accuracy lines are nearly parallel to the x-axis and only minor changes can be seen in *weather* and *digital* type within range less than 0.5%.

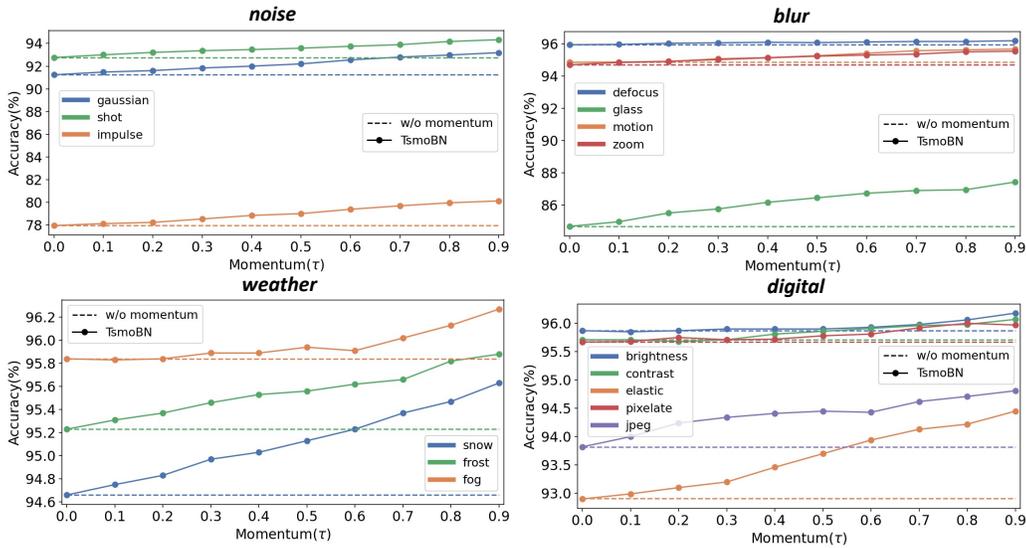


Figure E.4: Studies on how momentum values would affect TsmoBN performance. We report FedAvg on four unseen corruption types (i.e., noise, blur, weather and digital) in CIFAR-10-C.

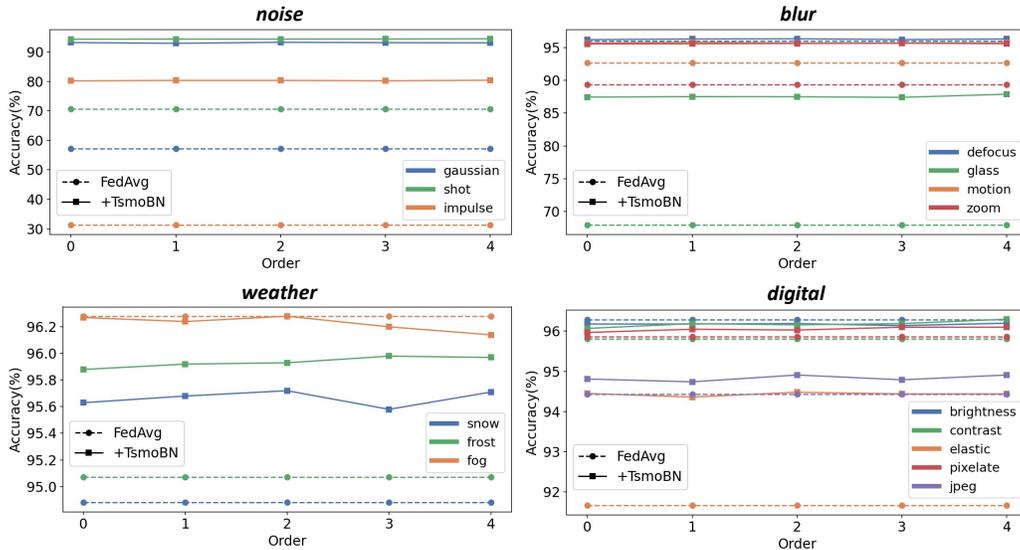


Figure E.5: Studies show batch sampling order would not affect TsmoBN. We report results of FedAvg and TsmoBN on CIFAR-10-C.

E.2.5 MORE EXPERIMENTAL RESULTS ON EXTRA FOUR UNSEEN CLIENTS FROM CIFAR-10-C

Besides the leave-one-CorruptionType-out experiments on CIFAR-10-C, we use four extra corruption types provided by CIFAR-10-C to conduct more evaluations, the samples of these four corruption types are shown in Fig. E.6. Specifically, we collect four models from leave-one-CorruptionType-out validation and report all results on the extra four unseen clients, as shown in Table E.11, Table E.12 and Table E.13. The 'Leaved CorruptionType' indicates that models do not see a specific type of corruptions at all. Our TsmoBN achieves consistent improvements with different FL algorithms or domain generalization methods on all 4 unseen clients averagely. From Table E.11, we can see that if the model has not seen any types of *noise*, it performs worse on *speckle noise* than the model (e.g., the row of *blur*) which has seen *noise*. For *gaussian blur* and *spatter*, since they do not strongly change



Figure E.6: Samples of four extra corruption types (i.e., *speckle noise*, *gaussian blur*, *spatter* and *saturate*) in CIFAR-10-C.

distributions, the results between different models are similar. As the *saturate* is totally different from all types in *digital* group, we can see the results in the *digital* row of Table E.12 are not much lower than other rows even if the model has not seen *digital* types of corruptions. Hence, our results show that involving more different distributions at training time improves the model generalizability on new distributions that does not change in a large degree, but for the very different new distribution, the performance is still degraded.

Leaved CorruptionType	Method	<i>speckle noise</i>	<i>gaussian blur</i>	<i>spatter</i>	<i>saturate</i>
<i>noise</i>	FedAvg	71.96(0.21)	96.53(0.13)	95.13(0.16)	92.16(0.58)
	+TsmoBN +6.44	93.68(0.36)	96.37(0.04)	95.87(0.14)	95.64(0.38)
	FedProx	72.68(2.53)	96.49(0.14)	94.88(0.48)	91.99(0.86)
	+TsmoBN +6.19	93.27(0.59)	96.33(0.11)	95.74(0.28)	95.45(0.44)
	FedNova	85.39(0.33)	96.91(0.10)	96.35(0.08)	93.74(0.47)
+TsmoBN +3.32	95.94(0.16)	96.86(0.04)	96.62(0.02)	96.25(0.16)	
FedAdam	72.44(1.92)	96.57(0.18)	95.23(0.06)	92.81(0.04)	
+TsmoBN +6.27	93.83(0.15)	96.48(0.09)	96.03(0.01)	95.81(0.20)	
FedBN	70.32(3.76)	96.23(0.23)	93.96(0.87)	90.44(1.67)	
+TsmoBN +6.92	92.32(1.18)	96.23(0.17)	95.34(0.39)	94.76(0.58)	
<i>blur</i>	FedAvg	89.15(1.75)	96.17(0.16)	95.83(0.19)	91.41(0.72)
	+TsmoBN +2.78	95.52(0.18)	96.28(0.10)	96.12(0.12)	95.77(0.15)
	FedProx	87.28(1.77)	96.21(0.12)	95.63(0.19)	90.39(1.39)
	+TsmoBN +3.50	95.34(0.15)	96.34(0.22)	96.15(0.02)	95.67(0.18)
	FedNova	95.77(0.39)	96.69(0.05)	96.57(0.10)	93.41(0.55)
+TsmoBN +0.88	96.36(0.18)	96.75(0.03)	96.62(0.09)	96.23(0.15)	
FedAdam	90.60(2.21)	96.37(0.12)	96.03(0.12)	91.95(0.34)	
+TsmoBN +2.41	95.83(0.20)	96.46(0.14)	96.42(0.13)	95.87(0.22)	
FedBN	86.50(1.81)	96.04(0.18)	95.42(0.38)	89.13(1.09)	
+TsmoBN +3.93	95.06(0.17)	96.39(0.14)	95.99(0.16)	95.39(0.47)	

Table E.11: Performance of TsmoBN for unseen client generalization on CIFAR-10-C dataset. We use models from leave-one-CorruptionType-out validation (i.e., leave-*noise*-out and leave-*blur*-out) and test the models on 4 extra unseen clients. We implement our approach on five state-of-the-art FL algorithms and report standard deviation (in bracket) for multiple runs.

Leaved CorruptionType	Method	<i>speckle noise</i>	<i>gaussian blur</i>	<i>spatter</i>	<i>saturate</i>
<i>weather</i>	FedAvg	88.37(0.91)	96.27(0.11)	95.70(0.11)	91.04(0.06)
	+TsmoBN +3.04	95.46(0.09)	96.32(0.16)	96.18(0.15)	95.58(0.18)
	FedProx	87.87(3.09)	96.29(0.10)	95.75(0.07)	91.03(0.66)
	+TsmoBN +3.17	95.52(0.16)	96.35(0.10)	96.14(0.12)	95.63(0.09)
	FedNova	95.11(0.64)	96.73(0.04)	96.47(0.07)	92.09(0.59)
	+TsmoBN +1.40	96.33(0.06)	96.79(0.09)	96.61(0.03)	96.25(0.15)
FedAdam	83.99(5.69)	95.82(0.95)	94.85(1.59)	88.62(4.28)	
+TsmoBN +4.52	94.54(1.64)	96.19(0.35)	95.69(0.89)	94.93(1.17)	
FedBN	84.91(1.61)	96.06(0.15)	95.30(0.29)	88.53(0.90)	
+TsmoBN +4.45	95.02(0.18)	96.34(0.15)	95.91(0.19)	95.35(0.24)	
<i>digital</i>	FedAvg	89.22(1.71)	96.00(0.02)	95.74(0.17)	89.91(0.87)
	+TsmoBN +3.15	95.52(0.07)	96.20(0.13)	96.11(0.11)	95.64(0.13)
	FedProx	88.53(2.31)	95.97(0.05)	95.66(0.17)	89.26(0.52)
	+TsmoBN +3.48	95.45(0.17)	96.26(0.13)	96.13(0.09)	95.49(0.12)
	FedNova	95.50(0.29)	96.59(0.13)	96.60(0.05)	92.56(0.12)
	+TsmoBN +1.13	96.31(0.02)	96.72(0.08)	96.64(0.06)	96.12(0.10)
FedAdam	88.23(4.75)	96.19(0.16)	95.84(0.29)	90.32(1.68)	
+TsmoBN +3.32	95.50(0.45)	96.45(0.14)	96.22(0.08)	95.67(0.37)	
FedBN	84.68(3.18)	95.52(0.55)	94.95(0.78)	86.50(1.76)	
+TsmoBN +5.05	94.71(0.61)	96.26(0.08)	95.91(0.34)	94.96(0.60)	

Table E.12: Continue from Table E.11. Performance of TsmoBN for unseen client generalization on CIFAR-10-C dataset. We use models from leave-one-CorruptionType-out validation (i.e., leave-*weather*-out and leave-*digital*-out) and test the models on 4 extra unseen clients. We implement our approach on five state-of-the-art FL algorithms and report standard deviation (in bracket) for multiple runs.

Leaved CorruptionType	Method	<i>speckle noise</i>	<i>gaussian blur</i>	<i>spatter</i>	<i>saturate</i>
<i>noise</i>	JiGen	70.87(0.62)	96.53(0.14)	94.12(0.59)	89.04(1.02)
	+TsmoBN +4.37	86.57(1.02)	95.40(0.33)	93.44(0.42)	92.63(0.55)
	CSD	78.38(0.29)	96.61(0.04)	95.70(0.08)	92.62(0.52)
	+TsmoBN +5.05	94.67(0.33)	96.68(0.05)	96.34(0.10)	95.81(0.07)
	RandConv	69.14(3.50)	79.74(2.90)	72.70(3.43)	74.09(3.06)
	+TsmoBN +18.36	89.43(1.54)	94.26(0.91)	92.26(1.05)	93.15(1.09)
<i>blur</i>	JiGen	82.24(5.41)	95.25(2.09)	94.35(2.40)	86.24(3.20)
	+TsmoBN +3.05	90.17(3.63)	94.58(2.58)	93.25(3.01)	92.26(3.29)
	CSD	92.58(1.40)	96.36(0.07)	96.16(0.23)	91.79(0.71)
	+TsmoBN +1.99	95.64(0.11)	96.52(0.09)	96.38(0.07)	96.03(0.18)
	RandConv	73.19(5.50)	77.19(0.84)	73.66(0.84)	71.02(1.28)
	+TsmoBN +16.13	88.98(0.93)	92.10(0.97)	90.70(0.93)	87.80(6.36)
<i>weather</i>	JiGen	86.03(0.20)	96.47(0.08)	95.50(0.18)	88.08(0.96)
	+TsmoBN +2.94	92.89(0.30)	95.99(0.04)	95.01(0.33)	93.94(0.14)
	CSD	92.88(1.19)	96.41(0.12)	96.03(0.16)	91.57(0.48)
	+TsmoBN +1.95	95.87(0.06)	96.51(0.01)	96.32(0.13)	95.99(0.05)
	RandConv	72.06(4.77)	78.96(2.87)	73.76(2.86)	73.02(2.95)
	+TsmoBN +17.55	90.34(1.63)	93.44(0.99)	91.92(1.25)	92.28(1.34)
<i>digital</i>	JiGen	79.50(1.55)	94.26(0.87)	93.55(0.86)	83.22(0.67)
	+TsmoBN +3.57	88.55(1.66)	93.46(0.93)	92.12(1.34)	90.66(1.19)
	CSD	91.81(0.74)	96.14(0.23)	95.76(0.30)	89.49(0.90)
	+TsmoBN +2.75	95.79(0.11)	96.39(0.06)	96.27(0.11)	95.78(0.13)
	RandConv	70.39(3.10)	74.98(2.49)	71.87(2.95)	68.79(2.96)
	+TsmoBN +18.34	88.27(0.14)	91.32(0.16)	89.91(0.17)	89.87(0.04)

Table E.13: Results of adding our proposed TsmoBN to current domain generalization methods on CIFAR-10-C dataset. We use models from leave-one-CorruptionType-out validation (i.e., leave-*noise*-out, leave-*blur*-out, leave-*weather*-out and leave-*digital*-out) and test the 4 models on 4 extra unseen clients. We report standard deviation (in bracket) for multiple runs.