

Birds of a Feather Help: Context-aware Client Selection for Federated Learning

Hangrui Cao^{1*}, Qiyang Pan^{1*}, Yifei Zhu¹, Jiangchuan Liu²

¹ Shanghai Jiao Tong University, Shanghai, China

² Simon Fraser University, Burnaby, Canada

hangrui@umich.edu, sim_arity@sjtu.edu.cn, yifei.zhu@sjtu.edu.cn, jcliu@sfu.ca

Abstract

Federated learning emerges as an effective technique to train neural network models collaboratively among individual edge devices without centralizing the raw data. However, the data heterogeneity in such a distributed system often breaks the classical IID assumption on data statistics and greatly harms the accuracy and convergence speed. In this paper, we propose to intelligently select clients through exploiting the data correlations among clients to improve learning performance. We propose a novel Neural Contextual Combinatorial Bandit approach, NCCB, that gracefully handles the non-trivial relationship between the extracted features and rewards, and satisfies the combinatorial constraints imposed by the federated learning. Extensive experiments on real-world datasets further demonstrate that our approach can achieve up to 50% higher speed to reach target accuracy, 17% higher accuracy than state-of-the-art solutions, and close-to-optimal performance.

1 Introduction

With the increasing awareness of privacy protection, laws and regulations, such as the GDPR in European Union (European Commission 2018) and the CCPA in California (Bukaty 2019), are established worldwide to restrict the access of raw edge data. This greatly challenges the traditional gather-and-analyze paradigm in data analysis, where a centralized server collects data from all clients into one place to conduct data-centric tasks. Federated learning is proposed to conduct neural network training via collaboration among clients and servers without uploading the raw data. Since its introduction, it has been widely studied and applied in various machine learning applications, including computer vision (Liu et al. 2020b), natural language processing (Yang et al. 2018), medical care (Huang et al. 2019) and etc.

While existing machine learning theories heavily rely on the assumption of identical independent distribution (IID) of raw data, clients' local data in the federated setting naturally is non-IID. This is also known as data heterogeneity, one of the unique challenges emerged from federated learning, along with device heterogeneity, network heterogeneity,

etc. The existence of data heterogeneity has severe consequences on various aspects of federated learning, including sub-optimal prediction performance, longer convergence time, and unfairness among clients (Kairouz et al. 2019).

Previous studies either ignore this non-IID challenge or assume some known local data statistics. For example, many local weight update algorithms are proposed still following the IID setting (Liu et al. 2020a; Uddin et al. 2020). In addition to these, Wang et al. propose a control algorithm to fully utilize all limited network resources under the assumption that datasets follow four specific distributions (Wang et al. 2019). Yu et al. propose to speed up convergence by assuming that the datasets follow identical distributions but may be correlated with each other (Yu, Yang, and Zhu 2019). More recent works start to leverage the posterior information, like local model weights (Wang et al. 2020; Chen et al. 2020), revealed during the training process to help future learning.

However, these attempts in handling the data heterogeneity challenge either fail to model the real world distributions accurately or incur large overhead in extracting posterior information. In this paper, we argue that datasets at different devices, though no longer follow the IID property, actually may correlate with each other, which can then be leveraged to help the federated learning process. For example, the corpus difference of two sport journalist can be smaller than that of one sport journalist and one literature professor, which may help the natural language processing training. The images captured by two cameras in similar scenarios are more likely to be similar than the images captured by the cameras that are deployed for different purposes, which may help the computer vision model training. In fact, such kind of correlations among different entities have been confirmed or validated in various disciplines. One representative example is the language usage correlation between different clients studied in sociology, where clients living in the same social circle also tend to use similar words or phrases, which is formally known as “language homophily”² (Aiello et al. 2012). Homophily can be vividly captured by an English proverb, “Birds of a feather flock together”, which also motivates our title (McPherson, Smith-Lovin, and Cook 2001). *Intuitively, if we could capture this type of correlation among individ-*

*These authors contributed equally. The corresponding author is Yifei Zhu.

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

²Homophily is generally believed to form due to the process of selection and social influence.

ual devices in federated learning, and consider such correlation when selecting clients, a good global model can be trained by more representative clients with fewer communication rounds.

However, multiple challenges directly emerge if we want to identify and leverage this type of correlation in federated model training. First, compared with traditional distributed machine learning, federated learning enforces tighter requirements on privacy, prohibiting the sharing of raw data. How to capture the correlation among local dataset without revealing private information remains a problem. Second, even if we have successfully extracted the features that lead to revelation of these correlations in a privacy-preserving way, the relationship between correlations and final performance is still unclear. Consequently, how to leverage this relationship to help the client selection process also needs to be designed. Third, rather than selecting one client, we need to select multiple clients per round in federated learning. Selecting m clients per training round from n total clients can have $O(n^m)$ choices in total. How to identify the best combinations in this exponential space is non-trivial.

In this paper, we present a context-aware client selection mechanism for federated learning based on a novel Neural Contextual Combinatorial Bandit (NCCB) approach. Our mechanism gracefully handles the privacy-aware context feature embedding, relationship learning, and combinatorial selection challenges. Specifically, our context-aware client selection mechanism consists of two parts, the context feature extraction part and the client selection part. In the first part, a local feature extractor is designed to embed the raw data vectors into a short-length feature vector via a light weight locality sensitive hashing function. In the client selection part, different clusters are first formed based on these feature vectors to preserve the correlation relationship among clients and reduce the profiling space. We then design a neural network to estimate the non-trivial relationship between context feature and reward. A bandit selector acts simultaneously to leverage the correlation among different clients and identify the right set of clients to maximize system utility in each round. In summary, our key contributions are as follows:

- We propose a context-aware client selection mechanism that cost-effectively exploits the correlation among clients and select the right set of clients to improve federated model training.
- We identify a light weight locality-sensitive hashing approach to represent user data with little overhead and privacy protected.
- A novel neural combinatorial contextual bandit approach is proposed to jointly learn the non-trivial context-reward relationship and satisfy the combinatorial requirement non-trivially.
- Extensive experiments on real-world datasets further validate the superior performances of our approach in accuracy and speed to reach the target accuracy.

2 System Model and Problem Formulation

In this section, we first model the overall federated learning system and clients' utility. Then we formulate our client selection problem as a combinatorial bandit problem.

System Model

In a typical federated learning system (McMahan et al. 2017), suppose there are n clients in total. For each client $i \in [1, 2, \dots, n]$, the corresponding dataset is D_i . A federated learning server selects m clients per round for model training with a total of T rounds. We denote these selected clients as B_t at round t . The server first sends the global weights to initialize the models in local devices³. After receiving the weights, each client trains the model based on its own dataset and obtains different weights matrix w_1, w_2, \dots, w_n . These local weights are then aggregated at the server side to update a new global model w_{global} , which are distributed to clients further to start a new iteration. The whole procedure is illustrated in Fig. 1(a). Let w^T denote the final model weight and $\sum_{i=1}^n l(w^T, D_i)$ measures the final loss of the model where l is the loss function. The whole federated learning process tries to minimize

$$\min \sum_{i=1}^n l(w^T, D_i) \quad (1)$$

We next model clients' utility in the federated settings considering two major factors. Training neural networks with a larger scale of dataset is more likely to result in a finer model that fits more variations of data. Dataset size of a client thus greatly affects its contribution to the model. In addition, a client with a larger training loss implies that the current global model performs worse locally and thus needs more training to improve the performance. Therefore, from the federated learning's perspective, the utility U of a client i with dataset D_i can be defined as

$$U(i) = |D_i| \sqrt{\frac{1}{|D_i|} \sum_{d \in D_i} Loss(d)^2} \quad (2)$$

where $Loss(d)$ is the loss function determined by local model. This utility form is also used in (Lai et al. 2021).

With more clients included in training, the global model performs better. Furthermore, since the global model is the aggregation of total m local models, as m increases, the contribution of adding one more client to the global model becomes weaker. Consequently, we define the utility of a combination of selected clients to be

$$U(B_t) = \sqrt{\sum_{i \in B_t} U(i)^2} = \sqrt{\sum_{i \in B_t} \left[|D_i| \sum_{d \in D_i} Loss(d)^2 \right]} \quad (3)$$

which is squared root of sum of squared utility functions of all training clients.

³Clients and devices are used interchangeably in this paper.

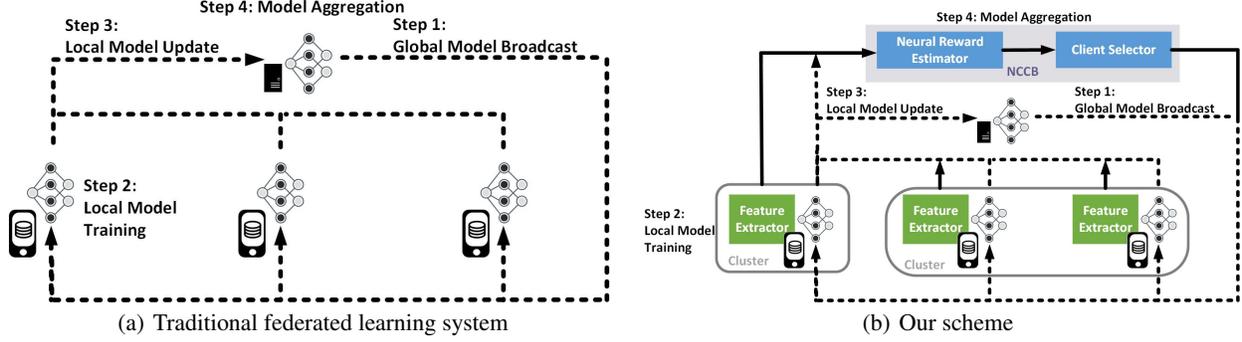


Figure 1: System overview

Problem Formulation

We study the client selection problem in federated learning so that the best representative clients can be selected to improve the federated learning process. Namely, a fixed number of clients has to be selected to join the federated learning in each round in order to maximize overall utility in Eq. 3.

In the federated learning process, since the utility of each user depends on its training loss, which further depends on the local data, user utilities remain unknown to the server. Exploration is required to train the model on a client and get the corresponding utility. As we explore more clients in the system, their utilities become gradually clear. An efficient system should also exploit previous experience in selecting multiple clients with large utilities. Therefore, the tradeoff of exploration and exploitation in the client selection process makes an appropriate formulation form.

We thus formulate our practical online client selection problem as a combinatorial bandit problem, in which the server and clients can be regarded as the player and arms. Note that unlike the classical bandit problem where only one arm needs to be selected, the combinatorial bandit here aims at capturing the practical training requirement in federated learning, where a set of clients, also known as an arm set (super arm) in bandit, has to be selected. We use a decision variable $B_t = (x_1^t, x_2^t, \dots, x_n^t), \forall t \in [1, 2, \dots, T]$ to denote the selection results in round t . If a client i is selected, its corresponding indication variable $x_i^t = 1$. Otherwise, $x_i^t = 0$. Denote the combination of optimal arms at round t is B_t^* . In the process of exploration and exploitation, the aim of combinatorial bandit is to select a set of arms in each round to minimize the cumulative regret. In summary, our client selection problem is formulated as

$$\begin{aligned} \min_x \quad & \mathbb{E} \left[\sum_{t=1}^T (U(B_t^*) - U(B_t)) \right] \\ \text{s.t.} \quad & \begin{cases} \sum_{i=1}^n x_i^t \leq m, \forall t \in [1, 2, \dots, T] \\ x_i^t \in [0, 1], \forall i \in [1, 2, \dots, n], t \in [1, 2, \dots, T] \end{cases} \end{aligned} \quad (4)$$

The first constraint limits the number of selected clients for each round. The second indicates the binary selection.

3 Design and Analysis

In this section, we first provide an overview for our context-aware client selection mechanism. We then introduce the details of our mechanism and conduct complexity analysis after that.

Overview

Our context-aware client selection consists of three parts, a local feature extractor, a reward estimator, and a client selector as illustrated in Fig. 1(b). At the client side, a local feature extractor maps the raw data vector to a feature vector without revealing private information. At the server side, a reward estimator learns the relationship between feature vectors and actual rewards and estimates the next-round rewards. Based on the estimator and other parameters, the client selector returns a client list to minimize the regret based on our proposed neural combinatorial contextual bandit approach.

Privacy-Preserving Context Extractor

The key requirement in the context extractor is to preserve data utility without uploading the raw data. Existing solutions use the intermediate weights generated after training on all local devices to heuristically guide client selection. However, this posterior information requires extra model training and thus incurs extra computation cost (Wang et al. 2020; Chen et al. 2020). In contrast to these heavy weight approaches, we identify a light weight approach to represent user data with little overhead, and show that embedding this into our proposed contextual bandit approach, the performance of federated learning can be greatly improved.

Local sensitive hashing (LSH) hashes similar items into one bucket (Gionis 1999) and is widely-used to generate fingerprints for documents (Zhao et al. 2020) and images (Zhou et al. 2021). For a vector $\mathbf{v} = (v_1, v_2, \dots, v_k)$ with size k , n indices $\{i_1, i_2, \dots, i_n\}$ from $\{1, 2, \dots, k\}$ are chosen randomly. Hence, a new vector $\mathbf{v}' = (v_{i_1}, v_{i_2}, \dots, v_{i_n}) = (v'_1, v'_2, \dots, v'_n)$ of n directions is formed. n different scalar hash functions $h_i : \mathbb{R}^n \rightarrow \mathbb{Z}$ defined by Eq. 5 are used to project \mathbf{v}' .

$$h_i(\mathbf{v}') = \lfloor \frac{v'_i - b_i}{\eta} \rfloor \quad (5)$$

where b_i is uniformly generated in $[0, \eta)$ and η is the quantization step. The generated vector $\mathbf{h}(\mathbf{v}) = (h_1(\mathbf{v}'), h_2(\mathbf{v}'), \dots, h_n(\mathbf{v}'))$ extracts the contents of vector \mathbf{v} and can be indexed with an integer $g(\mathbf{h}(\mathbf{v})) : \mathbb{Z}^k \rightarrow \mathbb{Z}$. For two near vectors \mathbf{v}, \mathbf{u} their corresponding hashes $g(\mathbf{v}), g(\mathbf{u})$ are similar. Therefore, LSH is a desirable lightweight candidate in distilling the prior information of local data without revealing raw data to server. Among all the LSH algorithms, SimHash(Manku, Jain, and Sarma 2007) is widely used to process large-scale documents(Jafari et al. 2021) and has also been used in some federated scenario, like cookie management (Google 2018). In the implementation, we also take advantage of SimHash algorithm to extract content features of datasets.

The Neural Combinatorial Contextual Bandit

Algorithm 1: NCCB

Input: Number of rounds: T ; number of total arms: n ; number of clusters: c ; number of selected arms per round: m ; exploration threshold k

- 1: Observe contextual vectors $\mathbf{x}_i, i \in \{1, 2, \dots, n\}$
- 2: Cluster all the arms into c clusters C_1, C_2, \dots, C_c
- 3: Initialize all the counters $Ct(C_i)$ to 0;
- 4: Randomly initialize θ_0
- 5: **for** $t \leftarrow 1, \dots, T$ **do**
- 6: Initialize the selected arm set $B_t = \phi$
- 7: $C_{un} = \{C_i | Ct(C_i) \leq k\}$.
- 8: $C_{sel} =$ Randomly select $\min\{m, |C_{un}|\}$ clusters from C_{un}
- 9: **for** cluster $C \in C_{sel}$ **do**
- 10: Randomly select an arm b_i from C
- 11: $B_t = B_t \cup b_i$
- 12: **end for**
- 13: **if** $|C_{un}| < m$ **then**
- 14: **for** $i \leftarrow 1, \dots, n$ **do**
- 15: Compute upper confidence bound U_i using $f(\mathbf{x}_{i,t}; \theta_{t-1})$
- 16: **end for**
- 17: **for** $i \leftarrow 1, \dots, m - |C_{un}|$ **do**
- 18: $b_i \leftarrow \operatorname{argmax}_{j \in \{1, \dots, n\} \wedge b_j \notin B_t} U_i$
- 19: $B_t = B_t \cup b_i$
- 20: **end for**
- 21: **end if**
- 22: Play b_i and observe $r_{b_i,t}$ for all $b_i \in B_t$
- 23: $\theta_t \leftarrow$ train Neural Network using $\{\mathbf{x}_{b_j}\}_{b_j \in B_t, i \in \{1, \dots, t\}}, \{r_{b_i,i}\}_{b_i \in B_t, i \in \{1, \dots, t\}}$, and θ_{t-1}
- 24: Update counter $Ct(C_i)$ for all $b_j \in B_t$ and $b_j \in C_i$
- 25: **end for**

Based on the represented features, we propose a neural combinatorial contextual bandit (NCCB) approach to embed the feature as context and select clients intelligently. The NCCB approach is made up of 3 steps: clustering, client selector, and neural reward estimator. Unlike the previous bandits that either just focus on capturing the non-linear relationship or solely focus on selecting a combination of arms,

our bandit jointly considers these two.

Clustering Before any trial, the feature vector for user i is denoted as $\mathbf{x}_i, i \in \{1, 2, \dots, n\}$. We cluster all the users into c clusters denoted as C_1, C_2, \dots, C_c using the feature vectors. A counter is set for each cluster to indicate the number of times users in this cluster are selected. An exploring threshold $k \in \mathbb{N}$ controls the number of selection for each cluster. This prevents the abuse of some specific arms and increases explorations especially at early stage of training as we will see in the following. The system initializes the counter of each client to zero.

Client selector A set $C_{un} = \{C_i\}$ includes all the unexplored clusters whose counters are less than or equal to k is formed. One arm will be selected per chosen cluster. If the number of clusters of these kinds exceeds the required number of arms m , they will be selected uniformly at random. It may happen that the number of clusters in C_{un} is smaller than m . Then, the system computes the upper confidence bound U_i for each arm and determines the top $m - |C_{un}|$ arms with the largest $m - |C_{un}|$ upper confidence bounds from a neural reward estimator. The rewards of all the clients are then predicted and the remaining $m - |C_{un}|$ arms are selected as shown as following:

$$b_{|C_{un}|+1}, b_{|C_{un}|+2}, \dots, b_m = (\operatorname{argsort}_{i \in \{1, \dots, k\} \wedge i \notin [b_1, \dots, b_{|C_{un}|}] U_i}[l : r]) \quad (6)$$

where $j \in 1, 2, \dots, n$, $\operatorname{argsort}$ is a function that takes the selected U_i array, sort them in an ascending order based on the upper confidence bound values and return n indexes. $l = n - (m - |C_{un}|) + 1$ and $r = n$. The arms selected from unexplored clusters C_{un} should be excluded since repeated selections of one arm are prohibited. This procedure is displayed from line 18-21 in Alg. 1. After selection, the system increases the counter of each cluster according to the selected clients.

Neural reward estimator In the mean time, we also need a reward estimator f that takes an arm and its context and estimate the corresponding reward. Unlike the classical linear contextual bandit where the estimation function is assumed to be linear, in our approach, a neural network is used to handle the complex relationship between features and rewards in our problem, which could be parameterized as

$$f(\mathbf{x}; \theta) = \sqrt{m} \mathbf{W}_L \sigma(\mathbf{W}_{L-1} \sigma(\dots \sigma(\mathbf{W}_1 \mathbf{x}))) \quad (7)$$

where L is the number of layer, $\mathbf{W}_l, l \in [1, \dots, L]$ are matrices whose elements are coefficients, σ is an activation function and $\theta = [\operatorname{vec}(\mathbf{W}_1), \operatorname{vec}(\mathbf{W}_2), \dots, \operatorname{vec}(\mathbf{W}_L)]$.

The loss function is defined as

$$L(\theta) = \sum_{t=1}^t \sum_{b \in B_t} \frac{(f(\mathbf{x}_b; \theta) - r_b(t))^2}{2} + w\lambda \frac{\|\theta - \theta^0\|_2^2}{2} \quad (8)$$

where w is the network width, λ is the regularization parameter and θ^0 is the initial parameter of the neural network model. For $l \in [1, \dots, L-1]$,

$$\mathbf{W}_l^0 = \begin{pmatrix} \mathbf{W} & \mathbf{0} \\ \mathbf{0} & \mathbf{W} \end{pmatrix} \quad (9)$$

where elements in \mathbf{W} of size $\frac{w}{2} \times \frac{w}{2}$ are drawn from $N(0, \frac{4}{w})$. We also initialize

$$\mathbf{W}_L^0 = (\mathbf{w}^T, -\mathbf{w}^T) \quad (10)$$

where entries in \mathbf{w} of size $\frac{w}{2}$ are drawn from $N(0, \frac{2}{w})$.

Generalizing the neural contextual bandit in (Zhou, Li, and Gu 2020), which only aims at selecting one arm, the upper bound in our combinatorial contextual bandit is calculated as

$$U = f(\mathbf{x}; \boldsymbol{\theta}) + \gamma \sqrt{\frac{\nabla f(\mathbf{x}; \boldsymbol{\theta})^T \mathbf{Z}^{-1} \nabla f(\mathbf{x}; \boldsymbol{\theta})}{w}} \quad (11)$$

where \mathbf{Z} is iteratively updated by

$$\mathbf{Z} = \mathbf{Z} + \frac{\nabla f(\mathbf{x}; \boldsymbol{\theta})^T \nabla f(\mathbf{x}; \boldsymbol{\theta})}{w} \quad (12)$$

after each training and γ is the scaling factor.

After selecting m arms, the corresponding actual rewards are observed sequentially. The neural network will be trained whenever a new reward is observed, with its parameter $\boldsymbol{\theta}$ updated. The reward of each arm $r_{b_t, t}$ is calculated from the training loss using Eq 2. The detailed algorithm design is shown in Alg. 1.

Complexity Analysis

For each training epoch, executing line 6-12 in Alg. 1 takes $\mathcal{O}(c + m)$. The worst case of each training epoch is that the number of unexplored clusters $|C_{un}|$ is smaller than the required number of selected arms m . In this case, the system computes $f(\mathbf{x}_{i,t}; \boldsymbol{\theta}_{t-1})$ for n times. Each time computing takes $\mathcal{O}(w^3 L)$. Thus, executing line 13-20 in Alg. 1 takes $\mathcal{O}(w^3 L + m \log n)$. The training of deep neural network takes $\mathcal{O}(w^3 L E T^2)$ where E is the epoch number set for training neural network in line 23 in Alg. 1. The factor T^2 is related to the number of dataset. Hence, the time complexity is $\mathcal{O}(w^3 m T^3 L E + w^3 n L T)$.

4 Evaluation

Experiment Setup

Dataset description We evaluate our scheme using Superuser dataset (Internet Archive 2021) and Yelp dataset (Yelp 2004). Superuser is an online forum where users can post and answer questions and comment others' questions and answers. The raw dataset includes information of posts, votes as well as comments, with their corresponding user ids. Yelp is one of the biggest review websites where users post their reviews after visiting restaurants and other places. The dataset includes the information of users and reviews, with their corresponding user ids.

The raw superuser dataset consists of 194085 users and the yelp dataset consists of 2189457 users. The dataset is highly skewed with most of users having too few comments to be trained on. Therefore, we remove the users with less than 150 comments or posts, and identify 761 valid users in Superuser dataset and 700 users in Yelp dataset to join federated learning. For both of the two dataset, we follow the method in (Hard et al. 2018) to remove stop words and

convert texts into sequences of length 10 to form (x, y) data pairs. Besides having different scales, these two dataset also represent two different data heterogeneity situations. We calculate the earth mover's distance (Zhao et al. 2018) of 5 randomly chosen clients' vocabulary distribution to the global distribution⁴. Yelp dataset is tested to be more divergent than the superuser dataset with EMD of 2.932 and 2.178, respectively. These two datasets help us better examine the performance of our algorithm in different data and scale settings.

Federated learning setting We use the post bodies of different users to model the local dataset at different devices. These posts are used to perform the next-word-prediction task, a popular NLP application in federated learning. We re-process the train dataset size for each client to be the same to exclude the influence of data size. We use the stacked LSTM model to perform word prediction as described in (Sundermeyer, Schlüter, and Ney 2012) with embedding size 200 and LSTM size 200. For Superuser dataset, the number of LSTM layer is 2 while for Yelp dataset the number of layer is increased to 5, since the latter dataset is bigger. Learning rate and learning rate per local are set as $\gamma = 0.002$ and $\gamma_d = 0.9993$. Local epoch is set to $E = 5$ and the number of clients participated in training each round is set to $N = 5$. The training is done with a 24G-memory NVIDIA RTX 3090.

Comparison approaches We compare our approach with four other approaches: random selection (McMahan et al. 2017), Oort (Lai et al. 2021), k-LinUCB (Li et al. 2010), ClusterFL (Ouyang et al. 2021). Random selection, is the default selection approach proposed in FedAvg, where clients are randomly selected to join the federated learning. Oort is a state-of-the-art stochastic bandit solution for client selection without considering contextual information. k-LinUCB is a representative contextual combinatorial bandit solution. It assumes a linear relationship between reward and feature vector, and selects k clients with top k expected rewards. ClusterFL is a state-of-the-art similarity-aware client selection algorithm based on the posterior weight information. After having all clients take one extra round of training, users are grouped into different clusters and are then randomly selected. In k-LinUCB and Oort, the reward is defined the same as in Eq. 2. We also present the performance of FedAvg in the ideal IID setting to demonstrate how close our method is to the ideal case.

Evaluation metrics We select Top-1 accuracy and speed to reach target accuracy as evaluation metrics. The Top-1 accuracy refers to the conventional accuracy where the predicted word must be the expected answer. We also evaluate the convergence speed for the training process, which is represented by the number of rounds to reach the target accuracy. The target accuracy is the maximum accuracy that the model obtains under the random selection method.

Overall Performance

We summarize the performance of the terminal accuracy and the rounds to target accuracy in Table 1 and Table 2, respectively. Compared with the baseline performance, our NCCB

⁴Datasets with larger EMD values has higher divergence.

approach improves the accuracy by 17.429% for Superuser dataset and by 12.267% for Yelp dataset. In addition, we find that K-LinUCB does not perform even as good as the classical stochastic bandit approach, revealing the non-trivial relationship between the context features and the rewards, and the consequences of wrong assumptions. NCCB also achieves 3% higher accuracy and uses 27% less rounds to reach target accuracy than ClusterFL in Superuser dataset, which indicates the efficiency of bandit. Comparison with the ideal IID case also demonstrates that NCCB can achieve close to optimal performance.

Table 1: Summary of terminal accuracy and relative improvement

| Dataset | Method | Accuracy (%) | Improvement (%) |
|-----------|-------------|--------------|-----------------|
| Superuser | Rand | 10.035 | 0.0 |
| | Oort | 11.372 | 13.323 |
| | k-LinUCB | 11.122 | 10.832 |
| | ClusterFL | 11.437 | 13.545 |
| | NCCB | 11.784 | 17.429 |
| | FedAVG(IID) | 12.050 | 20.080 |
| Yelp | Rand | 11.176 | 0.0 |
| | Oort | 12.116 | 8.410 |
| | k-LinUCB | 11.810 | 5.673 |
| | ClusterFL | 12.087 | 8.151 |
| | NCCB | 12.547 | 12.267 |
| | FedAVG(IID) | 12.800 | 14.531 |

Table 2: Summary of number of rounds to target accuracy and relative improvement

| Dataset | Method | Rounds to Target Accuracy | Improvement (%) |
|-----------|-------------|---------------------------|-----------------|
| Superuser | Rand | 283 | 0.0 |
| | Oort | 181 | 36.0 |
| | k-LinUCB | 213 | 24.7 |
| | ClusterFL | 184 | 34.90 |
| | NCCB | 134 | 52.7 |
| | FedAVG(IID) | 123 | 56.5 |
| Yelp | Rand | 196 | 0.0 |
| | Oort | 109 | 44.4 |
| | k-LinUCB | 108 | 44.9 |
| | ClusterFL | 114 | 41.8 |
| | NCCB | 101 | 48.5 |
| | FedAVG(IID) | 98 | 50.0 |

Sensitivity Analysis

Impact of exploring threshold The exploring threshold k is one important parameter in our method. It affects how we measure one cluster group as unexplored or not. We set k to be 0, 20, 100, 200 for the Superuser dataset, and 10, 15, 20, 200 for the Yelp dataset. Though NCCB achieves higher performances in all k settings than the widely applied random approach, different k leads to different performances for NCCB. For example, for the Yelp dataset, when we change k from 10 to 5, 20 and 200, the final accuracy decrease by 2.12%, 2.27% and 10.83% respectively. Essentially, k is a measurement of balance condition between exploration and exploitation. If k is very small (e.g. $k = 0$ in Fig. 2(a), $k = 5$ in Fig. 2(b)), it fails to explore enough

clients so that the neural reward estimator performs badly. On the other hand, if k is quite large (e.g. $k = 200$ in Fig. 2), the system neural reward estimator is idle during the whole process and fails to exploit previous experience.

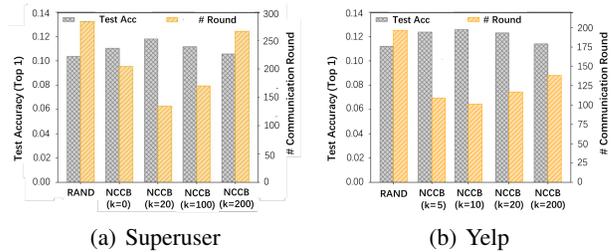


Figure 2: Impact of different exploring thresholds

Impact of scale of total participants We next explore the impact of total participated clients in training. As we can see, with the increase of number of participants, the training performance increases, indicating that NCCB method can capture clients with larger utility and improve the training performance. Further more, compared to the baseline with larger scale (761/700 training clients), our method can identify desirable clients from a smaller pool of clients and achieve even better performances.

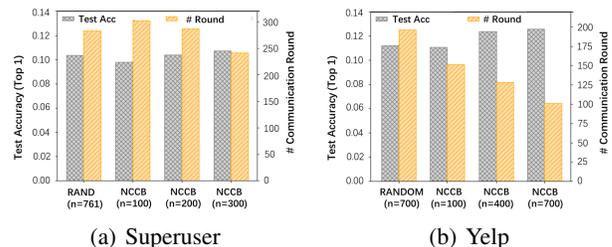


Figure 3: Impact of different system scales

Impact of feature size The feature size of the SimHash vector determines the expressiveness on clients' data. We evaluate the performance of NCCB under different feature sizes generated by SimHash. Figure 4 shows that the decrease of feature size degrades the performance of federated learning. For the superuser dataset, decreasing size of the feature vectors from 64 to 32, 16, 8 would decrease the final accuracy by 3.42%, 5.49% and 12.16% respectively. For the yelp dataset, when we decrease the size of feature vectors in a similar manner. The longer the bit length used to encode the local data, the more specific this feature captures the characteristics of this client, which helps the client selection process.

5 Related Works

Client selection in federated learning Client selection is an important controlling knob in federated learning to handle various types of heterogeneity. From the solutions' perspective, existing client selection studies generally fall into three

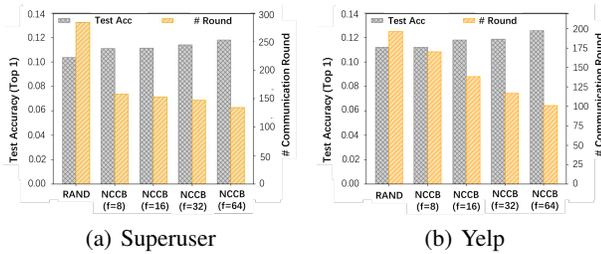


Figure 4: Impact of different feature sizes

categories: bandit-based (Xia et al. 2020; Lai et al. 2021), reinforcement learning based (Wang et al. 2020), heuristic based (Chen et al. 2020; Ouyang et al. 2021). Among bandit-based approaches, Oort in (Lai et al. 2021) formulates the client selection problem as a traditional multi-armed bandit problem without relying on contextual information. Consequently, it relies on training of all clients to get the quality estimation of local dataset, leading to prolonged convergence time. CS-UCB proposed in (Xia et al. 2020) relies on corresponding training time as the contextual feature to conduct UCB, which also requires redundant training rounds. In addition to bandit-based approaches, FAVOR in (Wang et al. 2020) applies a reinforcement learning approach to select device. However, deep reinforcement learning takes much longer time to train without any theoretical performance guarantee. Researchers also heuristically cluster clients based on the model weights after first-round training on all clients and randomly select clients without exploiting previous experience and incur significant cost (Chen et al. 2020; Ouyang et al. 2021).

Bandit algorithms Our work also contributes to the bandit field by jointly considering the non-trivial contextual relationship and the combinatorial requirement. Classical contextual bandit algorithms such as Linear Upper-Confidence-Bound (UCB) algorithm (Abe, Biermann, and Long 2003) assumes a trivial relationship between the explored rewards and feature vectors, which is hard to meet in real life (Zhou 2015). Neural contextual bandit is introduced to overcome this shortage by relying on a neural network to learn the underlying context-to-reward relationship (Zhou, Li, and Gu 2020). However, it fails to satisfy the combinatorial requirement specific in our problem and thus cannot be applied directly. Contextual combinatorial MAB (Chen, Xu, and Lu 2018) combines both contextual and combinatorial bandit. Nevertheless, it provides no method to estimate reward from feature vectors. In this paper, we propose to learn the non-trivial relationship between rewards and contexts via a neural network and satisfies the combinatorial requirements required in federated learning. The performance of NCCB is proved empirically.

6 Conclusion

In this paper, we proposed a novel context-aware client selection approach for communication-efficient federated learning by leveraging the correlation between clients. To be specific, we design a new neural contextual combinatorial

bandit approach to intelligently select clients to minimize the selection regret. Our approach successfully embeds the context into the selection process, handles the non-trivial relationship between the feature and reward, and works well in the intrinsic combinatorial searching space. Extensive empirical experiments on two large-scale real-world datasets further demonstrate its effectiveness.

References

- Abe, N.; Biermann, A. W.; and Long, P. M. 2003. Reinforcement learning with immediate rewards and linear hypotheses. *Algorithmica*, 37(4): 263–293.
- Aiello, L. M.; Barrat, A.; Schifanella, R.; Cattuto, C.; Markines, B.; and Menczer, F. 2012. Friendship prediction and homophily in social media. *ACM Trans. Web*, 6(2): 1–33.
- Bukaty, P. 2019. *The California Consumer Privacy Act (CCPA): An implementation guide*. IT Governance Publishing. ISBN 9781787781320.
- Chen, C.; Chen, Z.; Zhou, Y.; and Kailkhura, B. 2020. Fed-Cluster: Boosting the Convergence of Federated Learning via Cluster-Cycling. In *Proc. IEEE Int. Conf. Big Data*, 5017–5026.
- Chen, L.; Xu, J.; and Lu, Z. 2018. Contextual combinatorial multi-armed bandits with volatile arms and submodular reward. In *Proc. Adv. Neural Inf. Process. Syst.*, volume 31, 3247–3256.
- European Commission. 2018. 2018 reform of EU data protection rules. https://ec.europa.eu/commission/sites/beta-political/files/data-protection-factsheet-changes_en.pdf. Accessed: 2020-02-06.
- Gionis, A. 1999. Similarity Search in High Dimensions via Hashing. In *Proc. Int. Conf. Very Large Data Bases*.
- Google. 2018. Evaluation of cohort algorithms for the FLoC API. <https://github.com/google/ads-privacy/blob/master/proposals/FLoC/FLOC-Whitepaper-Google.pdf>. Accessed: 2020-02-06.
- Hard, A.; Rao, K.; Mathews, R.; Ramaswamy, S.; Beaufays, F.; Augenstein, S.; Eichner, H.; Kiddon, C.; and Ramage, D. 2018. Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*.
- Huang, L.; Shea, A. L.; Qian, H.; Masurkar, A.; Deng, H.; and Liu, D. 2019. Patient clustering improves efficiency of federated machine learning to predict mortality and hospital stay time using distributed electronic medical records. *J. Biomed. Inform.*, 99: 103291.
- Internet Archive. 2021. <https://archive.org/details/stackexchange>. Accessed: 2021-02-06.
- Jafari, O.; Maurya, P.; Nagarkar, P.; Islam, K. M.; and Crushev, C. 2021. A Survey on Locality Sensitive Hashing Algorithms and their Applications. *arXiv preprint arXiv:2102.08942*.
- Kairouz, P.; McMahan, H. B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A. N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. 2019. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*.

- Lai, F.; Zhu, X.; Madhyastha, H. V.; and Chowdhury, M. 2021. Oort: Efficient federated learning via guided participant selection. In *Proc. USENIX OSDI*, 19–35.
- Li, L.; Chu, W.; Langford, J.; and Schapire, R. E. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proc. ACM WWW*, 661–670.
- Liu, W.; Chen, L.; Chen, Y.; and Zhang, W. 2020a. Accelerating federated learning via momentum gradient descent. *IEEE Trans. Parallel Distrib. Syst.*, 31(8): 1754–1766.
- Liu, Y.; Huang, A.; Luo, Y.; Huang, H.; Liu, Y.; Chen, Y.; Feng, L.; Chen, T.; Yu, H.; and Yang, Q. 2020b. Fedvision: An online visual object detection platform powered by federated learning. In *Proc. AAAI Conf. Artif. Intell.*, volume 34, 13172–13179.
- Manku, G. S.; Jain, A.; and Sarma, A. D. 2007. Detecting Near-Duplicates for Web Crawling. In *Proc. ACM WWW*, 141–150.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Proc. Artif. Intell. Statist.*, 1273–1282.
- McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1): 415–444.
- Ouyang, X.; Xie, Z.; Zhou, J.; Huang, J.; and Xing, G. 2021. ClusterFL: a similarity-aware federated learning system for human activity recognition. In *Proc. ACM MobiSys*, 54–66.
- Sundermeyer, M.; Schlüter, R.; and Ney, H. 2012. LSTM neural networks for language modeling. In *Proc. Interspeech*.
- Uddin, M. P.; Xiang, Y.; Lu, X.; Yearwood, J.; and Gao, L. 2020. Mutual Information Driven Federated Learning. *IEEE Trans. Parallel Distrib. Syst.*, 32(7): 1526–1538.
- Wang, H.; Kaplan, Z.; Niu, D.; and Li, B. 2020. Optimizing federated learning on non-IID data with reinforcement learning. In *Proc. IEEE Conf. Comput. Commun.*, 1698–1707.
- Wang, S.; Tuor, T.; Salonidis, T.; Leung, K. K.; Makaya, C.; He, T.; and Chan, K. 2019. Adaptive federated learning in resource constrained edge computing systems. *IEEE J. Sel. Areas Commu.*, 37(6): 1205–1221.
- Xia, W.; Quek, T. Q. S.; Guo, K.; Wen, W.; Yang, H. H.; and Zhu, H. 2020. Multi-Armed Bandit-Based Client Scheduling for Federated Learning. *IEEE Trans. Wireless Commun.*, 19(11): 7108–7123.
- Yang, T.; Andrew, G.; Eichner, H.; Sun, H.; Li, W.; Kong, N.; Ramage, D.; and Beaufays, F. 2018. Applied federated learning: Improving google keyboard query suggestions. *arXiv preprint arXiv:1812.02903*.
- Yelp. 2004. Yelp open dataset. <https://www.yelp.com/dataset>. Accessed: 2020-02-06.
- Yu, H.; Yang, S.; and Zhu, S. 2019. Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning. In *Proc. AAAI Conf. Artif. Intell.*, volume 33, 5693–5700.
- Zhao, Y.; Li, M.; Lai, L.; Suda, N.; Civin, D.; and Chandra, V. 2018. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*.
- Zhao, Z.; Gao, M.; Luo, F.; Zhang, Y.; and Xiong, Q. 2020. LSHWE: improving similarity-based word embedding with locality sensitive hashing for cyberbullying detection. In *Int. Joint Conf. Neural Netw.*, 1–8.
- Zhou, D.; Li, L.; and Gu, Q. 2020. Neural contextual bandits with UCB-based exploration. In *Proc. Int. Conf. Mach. Learn.*, volume 119, 11492–11502.
- Zhou, J. T.; Zhang, L.; Jiawei, D.; Peng, X.; Fang, Z.; Xiao, Z.; and Zhu, H. 2021. Locality-Aware Crowd Counting. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Zhou, L. 2015. A survey on contextual multi-armed bandits. *arXiv preprint arXiv:1508.03326*.