INFOCOM'20

Optimizing Federated Learning on Non-IID Data with Reinforcement Learning

Hao Wang^{*}, Zakhary Kaplan^{*}, Di Niu[^], Baochun Li^{*} *University of Toronto, ^University of Alberta











Fedæcateel Learning



Federated Averaging Algorithm (FedAvg)

Initial model

Local data





Random selection

Local model

Local data







Thank you for the feedback

Local model









ML algorithms assume the training data is independent and identically distributed (IID)



algorithms but on non-ID data

Federated Learning reuses the existing ML







Non-ID data introduces bias into the and training failures

training and leads to a slow convergence

00000000000000000000 ススススタスススススススス 444444444444 MNIST 666666 フクフフフィアクククフラファフ 9999999999999999999

http://yann.lecun.com/exdb/mnist/



- FedAvg-IID





10 19 28 37 46 55 64 73 82 91 100 109 118 127 136 145 154 Communication Round (#)

Build IID training data?

No, we don't have any access to the data on your phone.





Zhao, Yue, et al. "Federated Learning with Non-IID Data." arXiv preprint arXiv:1806.00582 (2018).



Optimizing Federated Learning on Non-IID Data with Reinforcement Learning

[INFOCOM'20]

Build IID training data? No

data privacy





Peeking into the data distribution on each device without violating

Probing the bias of non-IID data



Carefully select devices to balance the bias introduced by non-IID data



Probing the data distribution



Non-IID data

Initial model

Local model

100 devices, each has 600 samples

666637 80% data has the same label, e.g, "6"

A two-layer CNN model with 431,080 parameters





We apply Principle Component Analysis (PCA) to reduce dimensionality







An implicit connection between model weights and data distribution





Probing the data distribution

Selecting devices for federated learning







K-Center Clustering

















Random Selection from Groups



















Communication Round (#)

Probing the data distribution

Selecting devices for federated learning

How to select devices to speed up training?



It is difficult to select the appropriate subset of devices - Model weights —> device selection choice - A dynamic and undeterministic problem

Reinforcement Learning (RL)





(..., state, action, reward, state', action', ..., end)



Episode





(..., state, action, rewar (..., state, action, rewar

(..., state, action, rewar
(..., state, action, rewar

d,	state',	action',	, en
d,	state',	action',	, en
d,	state',	action',	, en
d,	state',	action',	, en
d,	state',	action',	, en
d,	state',	action',	, en
• • •			
d,	state',	action',	, en
d,	state',	action',	,en



(..., state, action, reward, state', action', ..., end) (..., state, action, reward, state', action', ..., end) (..., state, action, reward, state', action', ..., end)

Learn to maximize sum(reward)

(..., state, action, reward, state', action', ..., end) (..., state, action, reward, state', action', ..., end)











Global weights Local model weights

100-dimension vector



Select K devices from a pool of N devices — a huge action space

Selecting 10 devices from a pool of 100 devices leads to **1.7310309e+13** possible actions

Actions

Modify the RL training algorithm

Selecting the Top K Devices

Only one device is selected during the RL training

devices from N devices

Now the action space is **{1, 2, ..., N}**, instead of selecting K

Evaluating Each Device







Rewards $r_t = \Xi(\omega_t - \Omega) - 1$ $0 \leq \omega_t \leq \Omega \leq 1$ $r_t \in (-1,0]$



	Positive constant
O_t	Training Accuracy
$\mathbf{\Omega}$	Target accuracy
t	Communication round #

More communication rounds: $t\uparrow$ —> sum(r_t) \downarrow

Training the DRL Agent

Look for a **function** that points out the **actions** leading to the maximum cumulative **return** under a particular **state**









Episode





Benchmark: MNIST, FashionMNIST, CIFAR-10 Non-IID level: 1, half-and-half, 80%, 50%

Half-and-half 3 3 3 3 3 7 7 7 7 7 7 66637 80%

Evaluating Our Solution





2200

Non-IID level



MNIST

FashionMNIST

CIFAR-10





Non-IID level half & half

Communication Rounds



MNIST

FashionMNIST

CIFAR-10





Non-IID level

80%







Non-IID level

50%



MNIST

FashionMNIST

CIFAR-10







Indirect data distribution probing DRL-based device selection Communication rounds can be reduced by up to 49% on the MNIST • 23% on FashionMNIST • 42% on CIFAR-10