

Demo code for FedAST: Federated Asynchronous Simultaneous Training

Used assets and codes

Our implementation is based on the public repositories of [FLSim](https://github.com/iQua/flsim) (<https://github.com/iQua/flsim>) and [Async-HFL](https://github.com/Orienfish/Async-HFL) (<https://github.com/Orienfish/Async-HFL>).

Requirements

We suggest creating a new environment before installing the required packages in order to avoid any collisions. We provide the versions of each package that we work with. They may work with other previously installed versions too.

To install requirements:

```
pip install -r requirements.txt
```

All experiments in the paper are run on NVIDIA GeForce GTX TITAN X at our internal cluster using CUDA. The code also supports running solely on the CPU!

Downloading Datasets

MNIST, Fashion-MNIST, CIFAR-10, and CIFAR-100 are available with built-in methods of PyTorch. They will be downloaded to `./data` automatically. We follow the instructions on [LEAF](https://github.com/TalwalkarLab/leaf) (<https://github.com/TalwalkarLab/leaf>) to download Shakespeare dataset. (You may use `--sf 0.1` or `--sf 0.2` for a subset of the dataset)

```
git clone https://github.com/TalwalkarLab/leaf.git
cd leaf/data/shakespeare/
./preprocess.sh -s niid --sf 1.0 -k 0 -t sample -tf 0.8
```

Running Demo Experiments

After installing the requirements, you can run the demo Notebook files. The config files we provide are set to use cuda as a default option. You need to change all occurrences of "cuda" to "cpu" in config files if CUDA is not available on your device. We tested both options on different devices.

demo_1 is a similar experiment to Figure 6.

You can also run any experiments on the terminal with the following line:

```
python run.py -c ./configFiles/**filename**.json
```

Configuration File Instructions

- Below, you can find the experimental parameters:

`num_of_threads` := Nb. of threads if only CPU is used (Not used in our experiments, you can set to 1), an integer value

`algo` := Training method, one of "fedast" (FedAST), "sync" (Sync-ST), "sync-bods" (Sync-Bayes-ST), "sync-ucb" (Sync-UCB-ST).

`partialNbofClients` := The total number of active clients, an integer value.

`full_async` := Flag to keep buffer size always at 1 (FedAST-NoBuffer), True for FedAST-NoBuffer.

`varBasedR` := Flag for dynamic / static options. True for dynamic, false for static.

`R_bs_ratio` := The value to keep the ratio of R/buffer_size fixed at, integer.

`max_time` := The maximum simulated time for the experiment, a float value.

`clients` := Client control settings

├—total := The number of total clients, an integer value.

├—slow_ratio := The ratio of slow clients, a float value between 0-1.

├—normal_ratio := The ratio of normal-speed clients, a float value between 0-1.

├—fast_ratio := The ratio of fast clients, a float value between 0-1.

`models` := Model control settings, sub-dictionaries should be named as "model1", "model2",...

├—model1 := Control settings of Model 1.

├—name := How many rounds Model 1 is trained, one of "MNIST", "FashionMNIST", "CIFAR-10", "Shakespeare", "CIFAR-100".

├—rounds := How many rounds Model 1 is trained, an integer value.

├—target_accuracy := Target test accuracy to train Model 1 until, a float value between 0-1.

├—model_path := "./models"

├—data_path := "./leaf/data/shakespeare/data" if name is "Shakespeare" else "./data"

├—model_device := Device to store model, one of cuda or cpu.

- |—batch_size:= Batch size, an integer value.
- |—buffer_size:= Buffer size for FedAST, an integer value.
- |—epochs:= If shuffled (epoch) training is done, a boolean value. If True, local_iter epochs are run for each local training, if False local_iter SGD steps are run on each local training.
- |—local_iter:= Number of SGD steps or epochs based on epochs value, an integer value.
- |—firstC:= k in accepted first-k updates for synchronous methods, an integer value.
- |—round_to_test:= Number of round period for test set evaluation, an integer value.
- |—local_lr:= Client-side learning rate, a float value.
- |—global_lr:= Server-side learning rate, a float value.
- |—nb_of_active_works:= Number of active local training requests for FedAST. The ratio of these values across models for synchronous methods determines client sharing. An integer value.
- |—delay:= Control settings of delay generation.
 - |—name:= shiftedExp.
 - |—exp_mean:= Mean of exponential part, a float value.
 - |—shift:= Constant shift part, a float value.
 - |—local_multi:= Number of local SGD steps-manuel entry. 27 for our experiments. An integer value.
- |—data:= Control settings of data.
 - |—trainset_device:= Device to store trainset, one of cuda or cpu.
 - |—testset_device:= Device to store testset, one of cuda or cpu.
 - |—data_seed:= Seed for random data distribution, an integer value.
 - |—dirichlet:= Dirichlet data distribution
 - |—alpha:= α parameter for Dirichlet data distribution
- |—model2 := Control settings of Model 2 (Same as Model 1).
- |—model3 := Control settings of Model 3 (Same as Model 1).
- |—model4 := Control settings of Model 4 (Same as Model 1).
- |—... := You can add as many models as desired.

bods_alpha := α parameter for Sync-Bayes-ST as used in the original paper, a float value between 0-1.

ucb_gamma := γ parameter for Sync-UCB-ST as used in the original paper, a float value between 0-1.