

Decentralized Learning Made Easy With DecentralizePy

Akash Dhasade

Anne-Marie Kermarrec

Rafael Pires

Rishi Sharma

Milos Vujasinovic

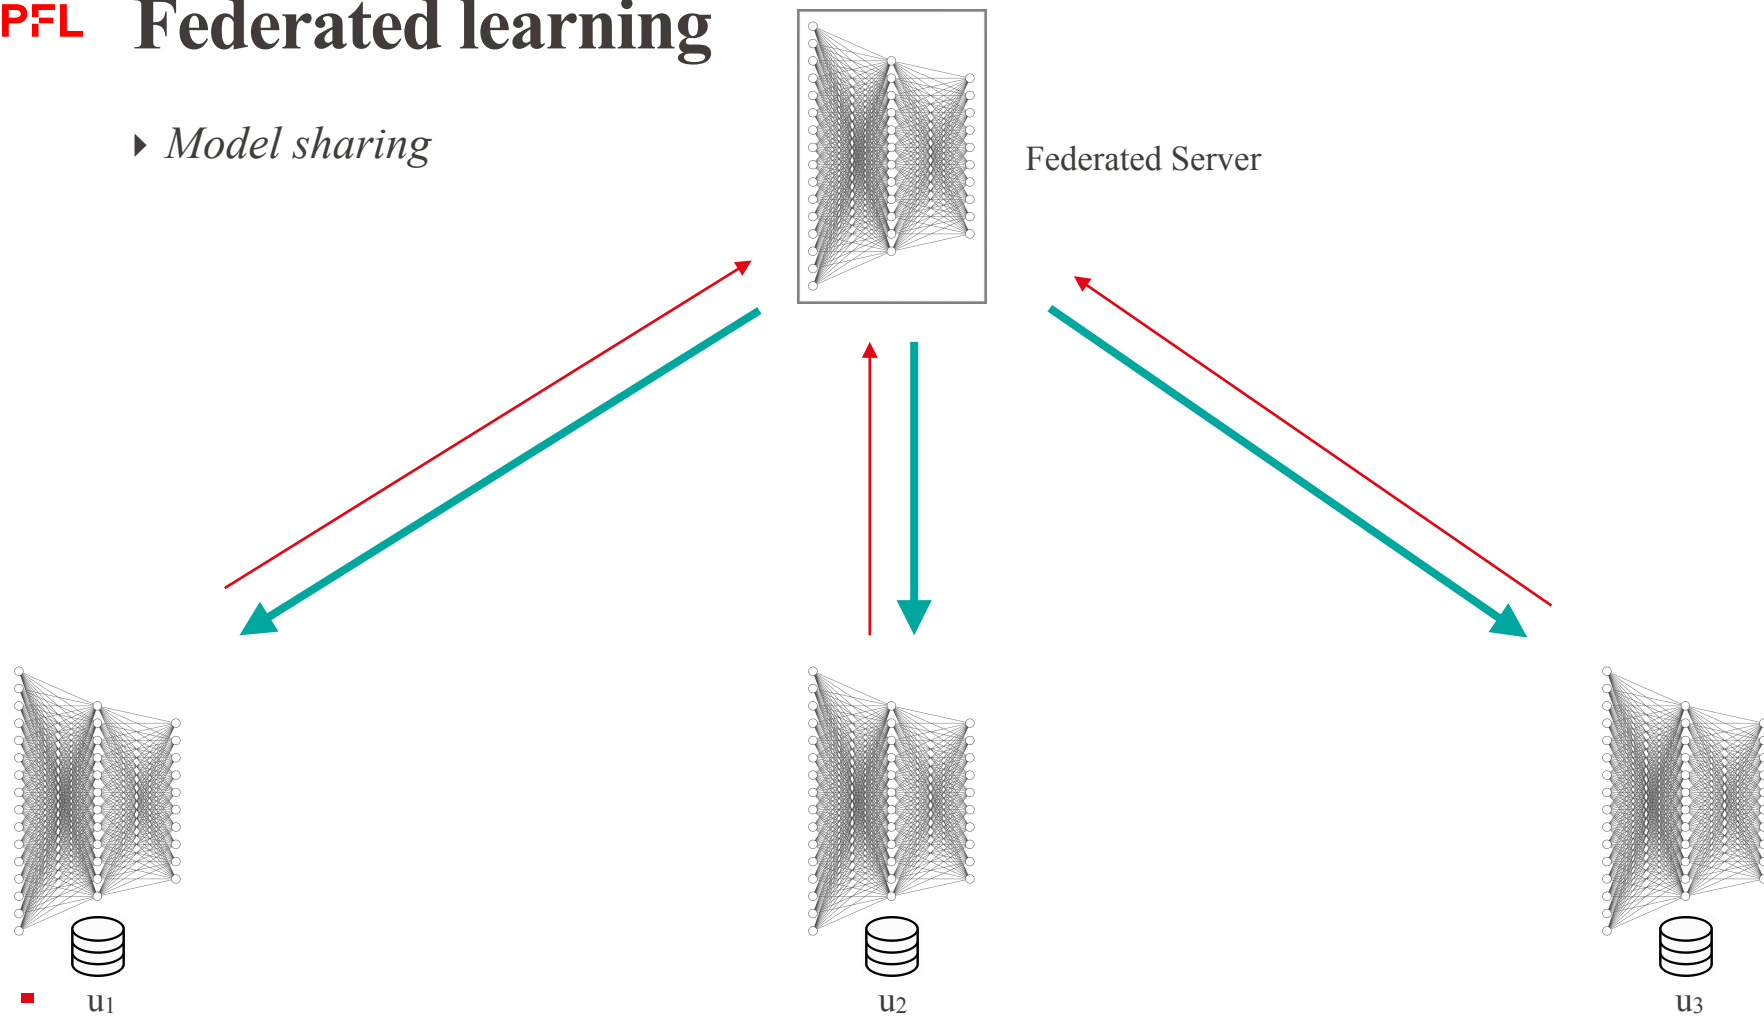
Scalable Computing Systems Laboratory
EPFL

A **framework** for designing and studying **decentralized learning systems**.

Rapid development

Scalability

► *Model sharing*



FedScale

A scalable and extensible federated learning engine and benchmark

GET STARTED

FedScale is a scalable and extensible **open-source** federated learning (FL) engine. It provides high-level APIs to implement FL algorithms, deploy and evaluate them at scale across diverse hardware and software backends. FedScale also includes the largest **FL benchmark** that contains FL tasks ranging from image classification and object detection to language modeling and speech recognition. Moreover, it includes datasets to faithfully emulate **FL runtime environments** where FL solutions will realistically be deployed.



We are actively developing FedScale, and welcome contributions from community. [Join our slack](#) to keep up to date.

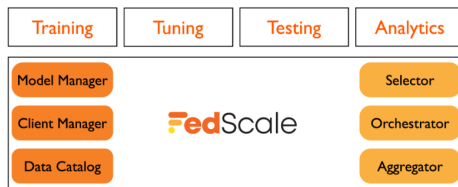
What's new? Flower Next Pilot Program >

Flower A Friendly Federated Learning Framework

A unified approach to federated learning, analytics, and evaluation. Federate any workload, any ML framework, and any programming language.

Take the tutorial

to learn federated learning



TensorFlow Federated: Machine Learning on Decentralized Data

TensorFlow Federated (TFF) is an open-source framework for machine learning and other computations on decentralized data. TFF has been developed to facilitate open research and experimentation with **Federated Learning (FL)**, an approach to machine learning where a shared global model is trained across many participating clients that keep their training data locally. For example, FL has been used to train **prediction models for mobile keyboards** without uploading sensitive typing data to servers.

TFF enables developers to simulate the included federated learning algorithms on their models and data, as well as to experiment with novel algorithms. Researchers will find **starting points and complete examples** for many kinds of research. The building blocks provided by TFF can also be used to implement non-learning computations, such as **federated analytics**. TFF's interfaces are organized in two main layers:

- Federated Learning (FL) API**
This layer offers a set of high-level interfaces that allow developers to apply the included implementations of federated training and evaluation to their existing TensorFlow models.
- Federated Core (FC) API**
At the core of the system is a set of lower-level interfaces for concisely expressing novel federated algorithms by combining TensorFlow with distributed communication operators within a strongly-typed functional programming environment. This layer also serves as the foundation upon which we've built Federated Learning.

TFF enables developers to declaratively express federated computations, so they could be deployed to diverse runtime environments. Included with TFF is a performant multi-machine simulation runtime for experiments. Please visit the [tutorials](#) and try it out yourself!

For questions and support, find us at the [tensorflow-federated](#) tag on StackOverflow.

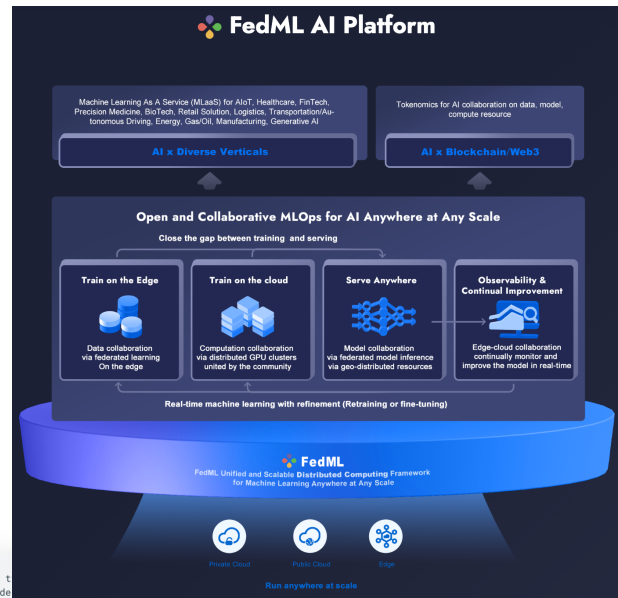
```
import tensorflow as tf
import tensorflow_federated as tff

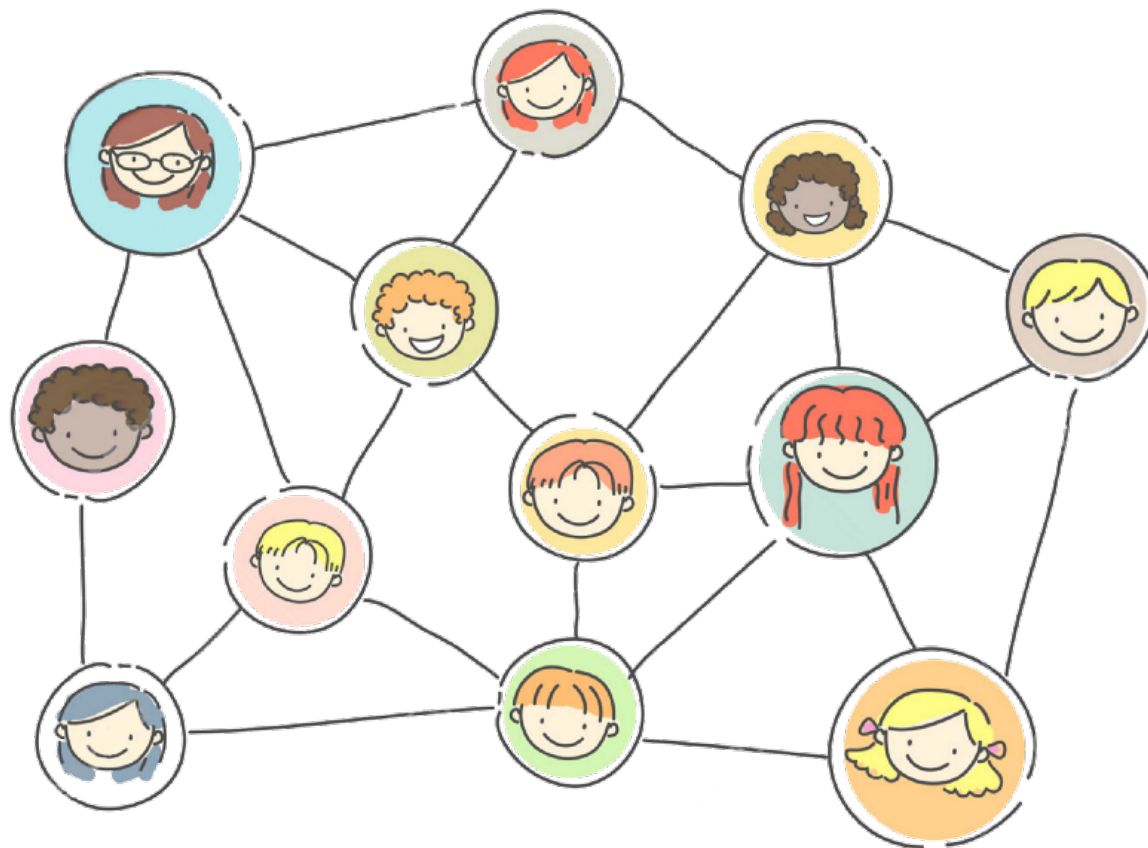
# Load simulation data.
source, _ = tff.simulation.datasets.emnist.load_data()
def client_data(n):
    return source.create_tf_dataset_for_client(source.client_ids[n]).map(
        lambda e: (tf.reshape(e['pixels'], [-1]), e['label'])
    ).repeat(10).batch(20)

# Pick a subset of client devices to participate in training.
train_data = [client_data(n) for n in range(3)]

# Wrap a Keras model for use with TFF.
def model_fn():
    model = tf.keras.models.Sequential([
        tf.keras.layers.Dense(10, tf.nn.softmax, input_shape=(784,),
            kernel_initializer='zeros')
    ])
    return tff.learning.models.from_keras_model(
        model,
        input_spec=tff.data.TensorSpec(0, element_spec,
            loss=tf.keras.losses.SparseCategoricalCrossentropy(),
            metrics=[tf.keras.metrics.SparseCategoricalAccuracy()])

# Simulate a few rounds of training with the selected client devices.
trainer = tff.learning.algorithms.build_weighted_fed_avg(
    model_fn,
    client_optimizer_fn=lambda: tf.keras.optimizers.SGD(0.1))
state = trainer.initialize()
for _ in range(5):
```





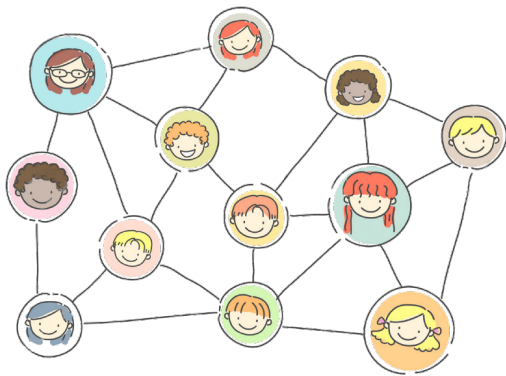
-

Train → Send



Receive → Aggregate

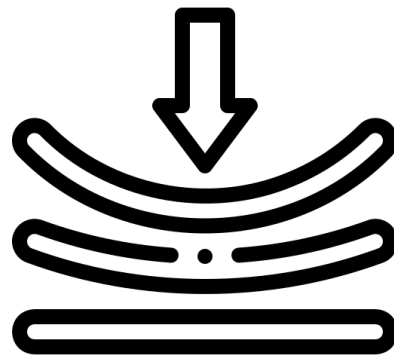
Simulations + No Re-usability!



Topology



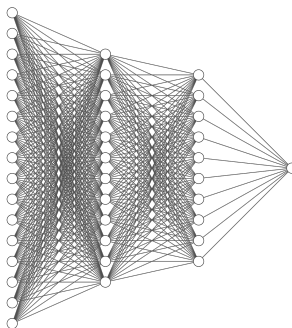
Communication



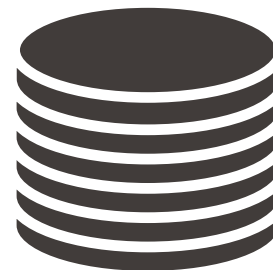
Compression



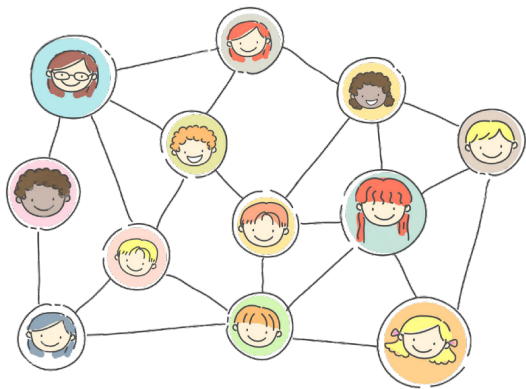
Roles



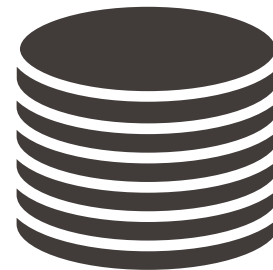
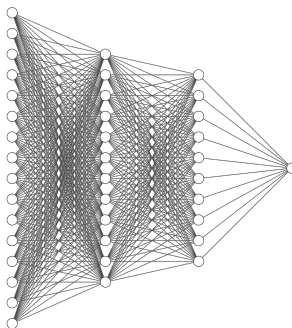
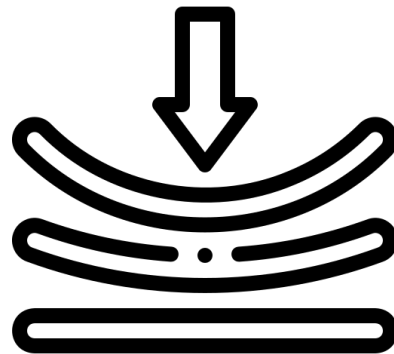
Models



Datasets



(Flexibility)




```
1 from decentralizepy.node.Node import Node
2
3 class DLNode(Node):
4     def run(self, iterations, training, dataset,
5             sharing, graph, communication):
6         for round in range(iterations):
7             training.train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send(neighbors, msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()
```

DecentralizePy already contains reference implementations of well-known algorithms.

```
1 from decentralizepy.node.Node import Node
2
3 class DLNode(Node):
4     def run(self, iterations, training, dataset,
5             sharing, graph, communication):
6         for round in range(iterations):
7             training.train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send(neighbors, msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()
```

```
1 from decentralizepy.node.Node import Node
2
3 class DLNode(Node):
4     def run(self, iterations, training, dataset,
5             sharing, graph, communication):
6         for round in range(iterations):
7             training.train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send(neighbors, msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()
```

```
1 from decentralizepy.node.Node import Node
2
3 class DLNode(Node):
4     def run(self, iterations, training, dataset,
5             sharing, graph, communication):
6         for round in range(iterations):
7             training.train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send(neighbors, msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()
```

```
1 from decentralizepy.node.Node import Node
2
3 class DLNode(Node):
4     def run(self, iterations, training, dataset,
5             sharing, graph, communication):
6         for round in range(iterations):
7             training.train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send(neighbors, msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()
```

```

1 from decentral_learner.node import Node
2
3 class DNode(Node):
4     def run(self, iterations, training_dataset,
5           sharing_graph, communication):
6         for round in range(iterations):
7             training_train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send_neighbors(msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()

```

```

1 from decentral_learner.node import Node
2
3 class DNode(Node):
4     def run(self, iterations, training_dataset,
5           sharing_graph, communication):
6         for round in range(iterations):
7             training_train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send_neighbors(msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()

```

```

1 from decentral_learner.node import Node
2
3 class DNode(Node):
4     def run(self, iterations, training_dataset,
5           sharing_graph, communication):
6         for round in range(iterations):
7             training_train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send_neighbors(msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()

```

```

1 from decentral_learner.node import Node
2
3 class DNode(Node):
4     def run(self, iterations, training_dataset,
5           sharing_graph, communication):
6         for round in range(iterations):
7             training_train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send_neighbors(msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()

```

```

1 from decentral_learner.node import Node
2
3 class DNode(Node):
4     def run(self, iterations, training_dataset,
5           sharing_graph, communication):
6         for round in range(iterations):
7             training_train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send_neighbors(msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()

```

```

1 from decentral_learner.node import Node
2
3 class DNode(Node):
4     def run(self, iterations, training_dataset,
5           sharing_graph, communication):
6         for round in range(iterations):
7             training_train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send_neighbors(msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()

```

```

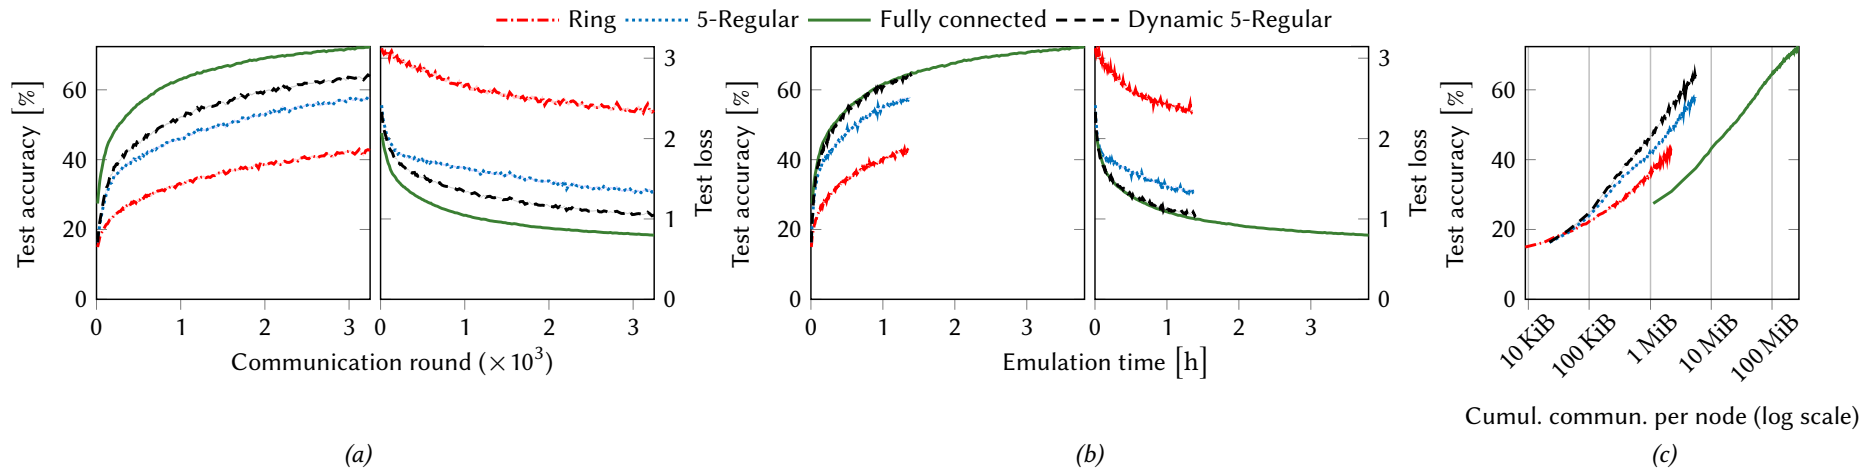
1 from decentral_learner.node import Node
2
3 class DNode(Node):
4     def run(self, iterations, training_dataset,
5           sharing_graph, communication):
6         for round in range(iterations):
7             training_train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send_neighbors(msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()

```

We use DecentralizePy as a **catalyst** for DL research in our lab.

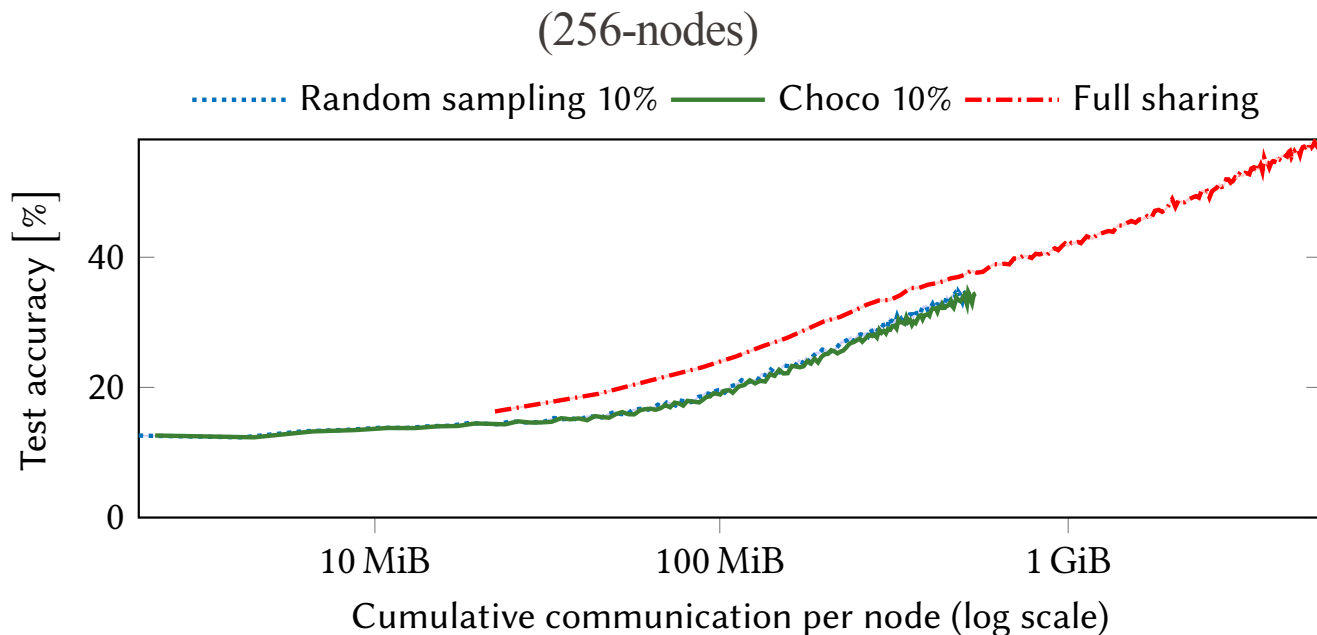
- ◆ CIFAR-10 (Non-IID) with GN-LeNet
- ◆ 256 and 1024 DL nodes
- ◆ Emulation on 16 machines
- ◆ D-PSGD with Metropolis Hastings

(256-nodes)



Information spreads faster through the network with dynamic topologies.

```
1 from decentralizepy.node.Node import Node
2
3 class DLNode(Node):
4     def run(self, iterations, training, dataset,
5             sharing, graph, communication):
6         for round in range(iterations):
7             training.train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send(neighbors, msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()
```



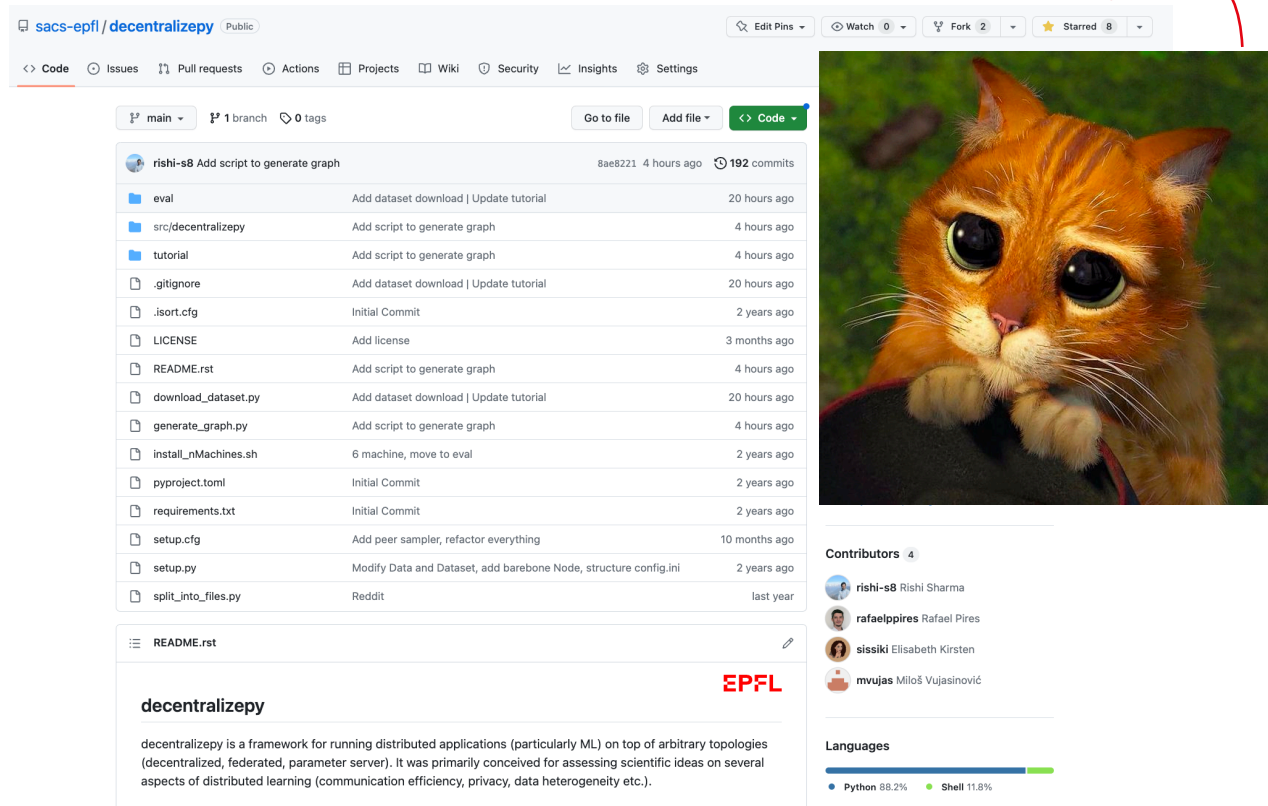
The loss of information due to compression dramatically affects the convergence in non-IID settings at scale.


```
1 from decentralizepy.node.Node import Node
2
3 class DLNode(Node):
4     def run(self, iterations, training, dataset,
5             sharing, graph, communication):
6         for round in range(iterations):
7             training.train(dataset)
8             msg = sharing.get_message()
9             neighbors = graph.get_neighbors()
10            communication.send(neighbors, msg)
11            rcv = communication.receive_from_all()
12            sharing.average(rcv)
13            dataset.test()
```

- ◆ Open source
- ◆ Already being used for a number of projects
- ◆ Adding new algorithms
- ◆ Realistic network emulations
- ◆ Peer-sampling and availability traces



<https://github.com/sacs-epfl/decentralizepy>



The screenshot displays the GitHub repository page for 'sacs-epfl/decentralizepy'. The repository is public and has 192 commits. The file structure includes folders for 'eval', 'src/decentralizepy', and 'tutorial', and files for '.gitignore', '.isort.cfg', 'LICENSE', 'README.rst', 'download_dataset.py', 'generate_graph.py', 'install_nMachines.sh', 'pyproject.toml', 'requirements.txt', 'setup.cfg', 'setup.py', and 'split_into_files.py'. The commit history shows recent updates to the dataset download script and the tutorial. The contributors list includes Rishi Sharma, Rafael Pires, Elisabeth Kirsten, and Miloš Vujanović. The languages section shows Python at 88.2% and Shell at 11.8%.

Contributors 4

- rishi-s8 Rishi Sharma
- rafaelpires Rafael Pires
- sisiki Elisabeth Kirsten
- mvujas Miloš Vujanović

Languages

- Python 88.2%
- Shell 11.8%

Please try DecentralizePy if you are working with DL and help us improve the framework.

Go Decentralized!



<https://github.com/sacs-epfl/decentralizepy>