

**ECML
PKDD
2024**



**UNIVERSITÀ
DI TORINO**

fl**u**k**e**: federated learning utility framework for experimentation and research

Mirko Polato, PhD

Department of Computer Science, University of Turin, Italy

mirko.polato@unito.it



2nd **W**orkshop on **A**dvancement in **F**ederated **L**earning @ ECML PKDD 2024

Yet another FL framework?


Not really!



from the idea



through the
implementation



to testing

Yet another FL framework?

Not really!


f1uke



from the idea



through the
implementation



to testing




simulated federation



deploy

Main features of fluke

Designed for fast prototyping and testing



<https://github.com/makgyver/fluke>

- **Open source:** fluke is an open-source Python package;
- **Easy to use:** fluke is designed to be extremely easy to use out of the box;
- **Easy to extend:** fluke is designed to minimize the overhead of adding new algorithms;
- **Up-to-date:** fluke comes with several (30+) state-of-the-art federated learning algorithms and datasets and it is regularly updated to include the latest affirmed techniques;
- **Easy to read:** the source code of the algorithms is written to mimic as close as possible the description in the reference papers.



flake is on PyPi!

Install it using a single command



```
$ pip install flake-fl
```

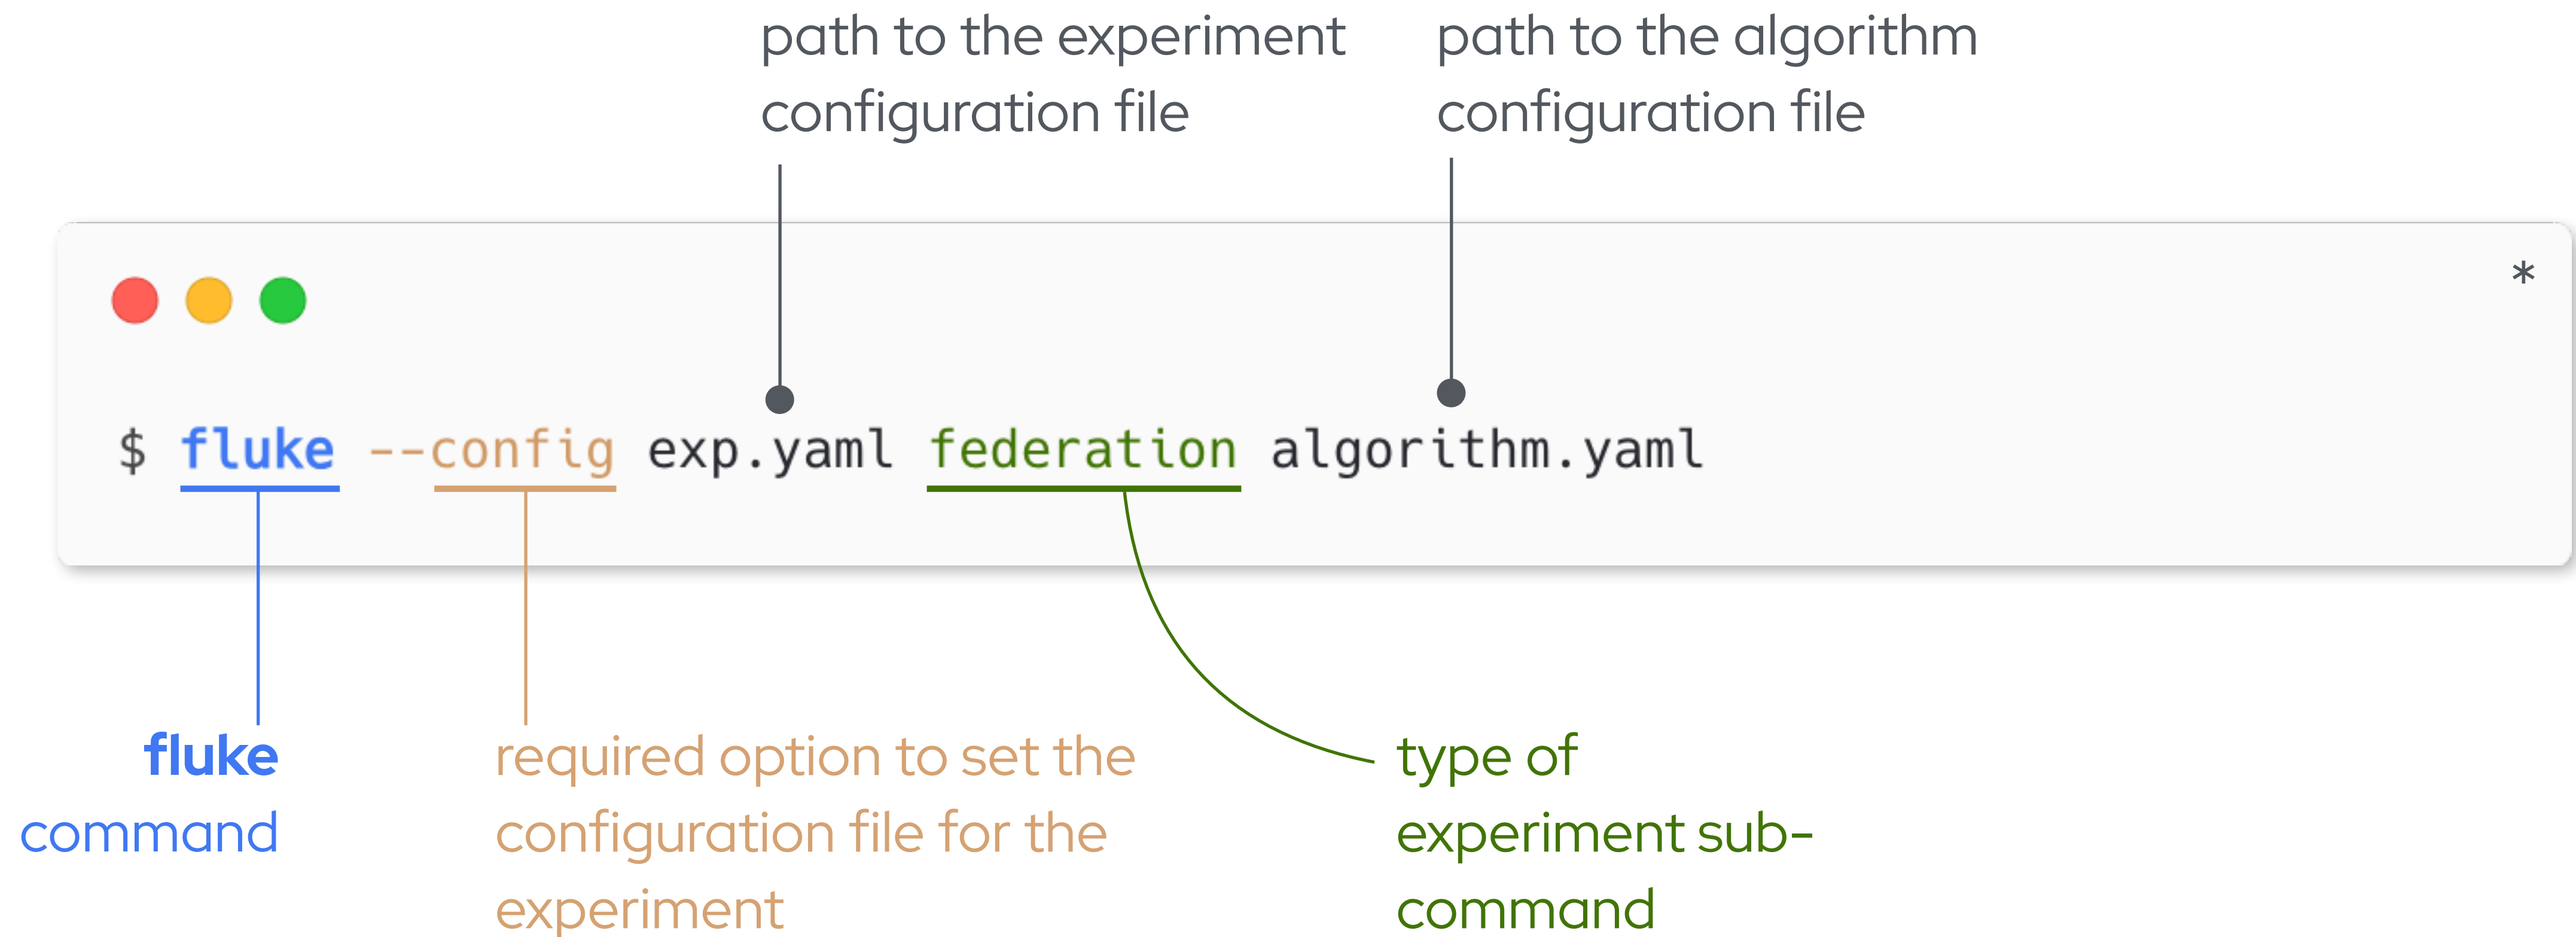
... or by cloning the repo



```
$ git clone https://github.com/makgyver/fluke.git  
$ cd fluke  
$ pip install -r requirements.txt
```

fluke CLI

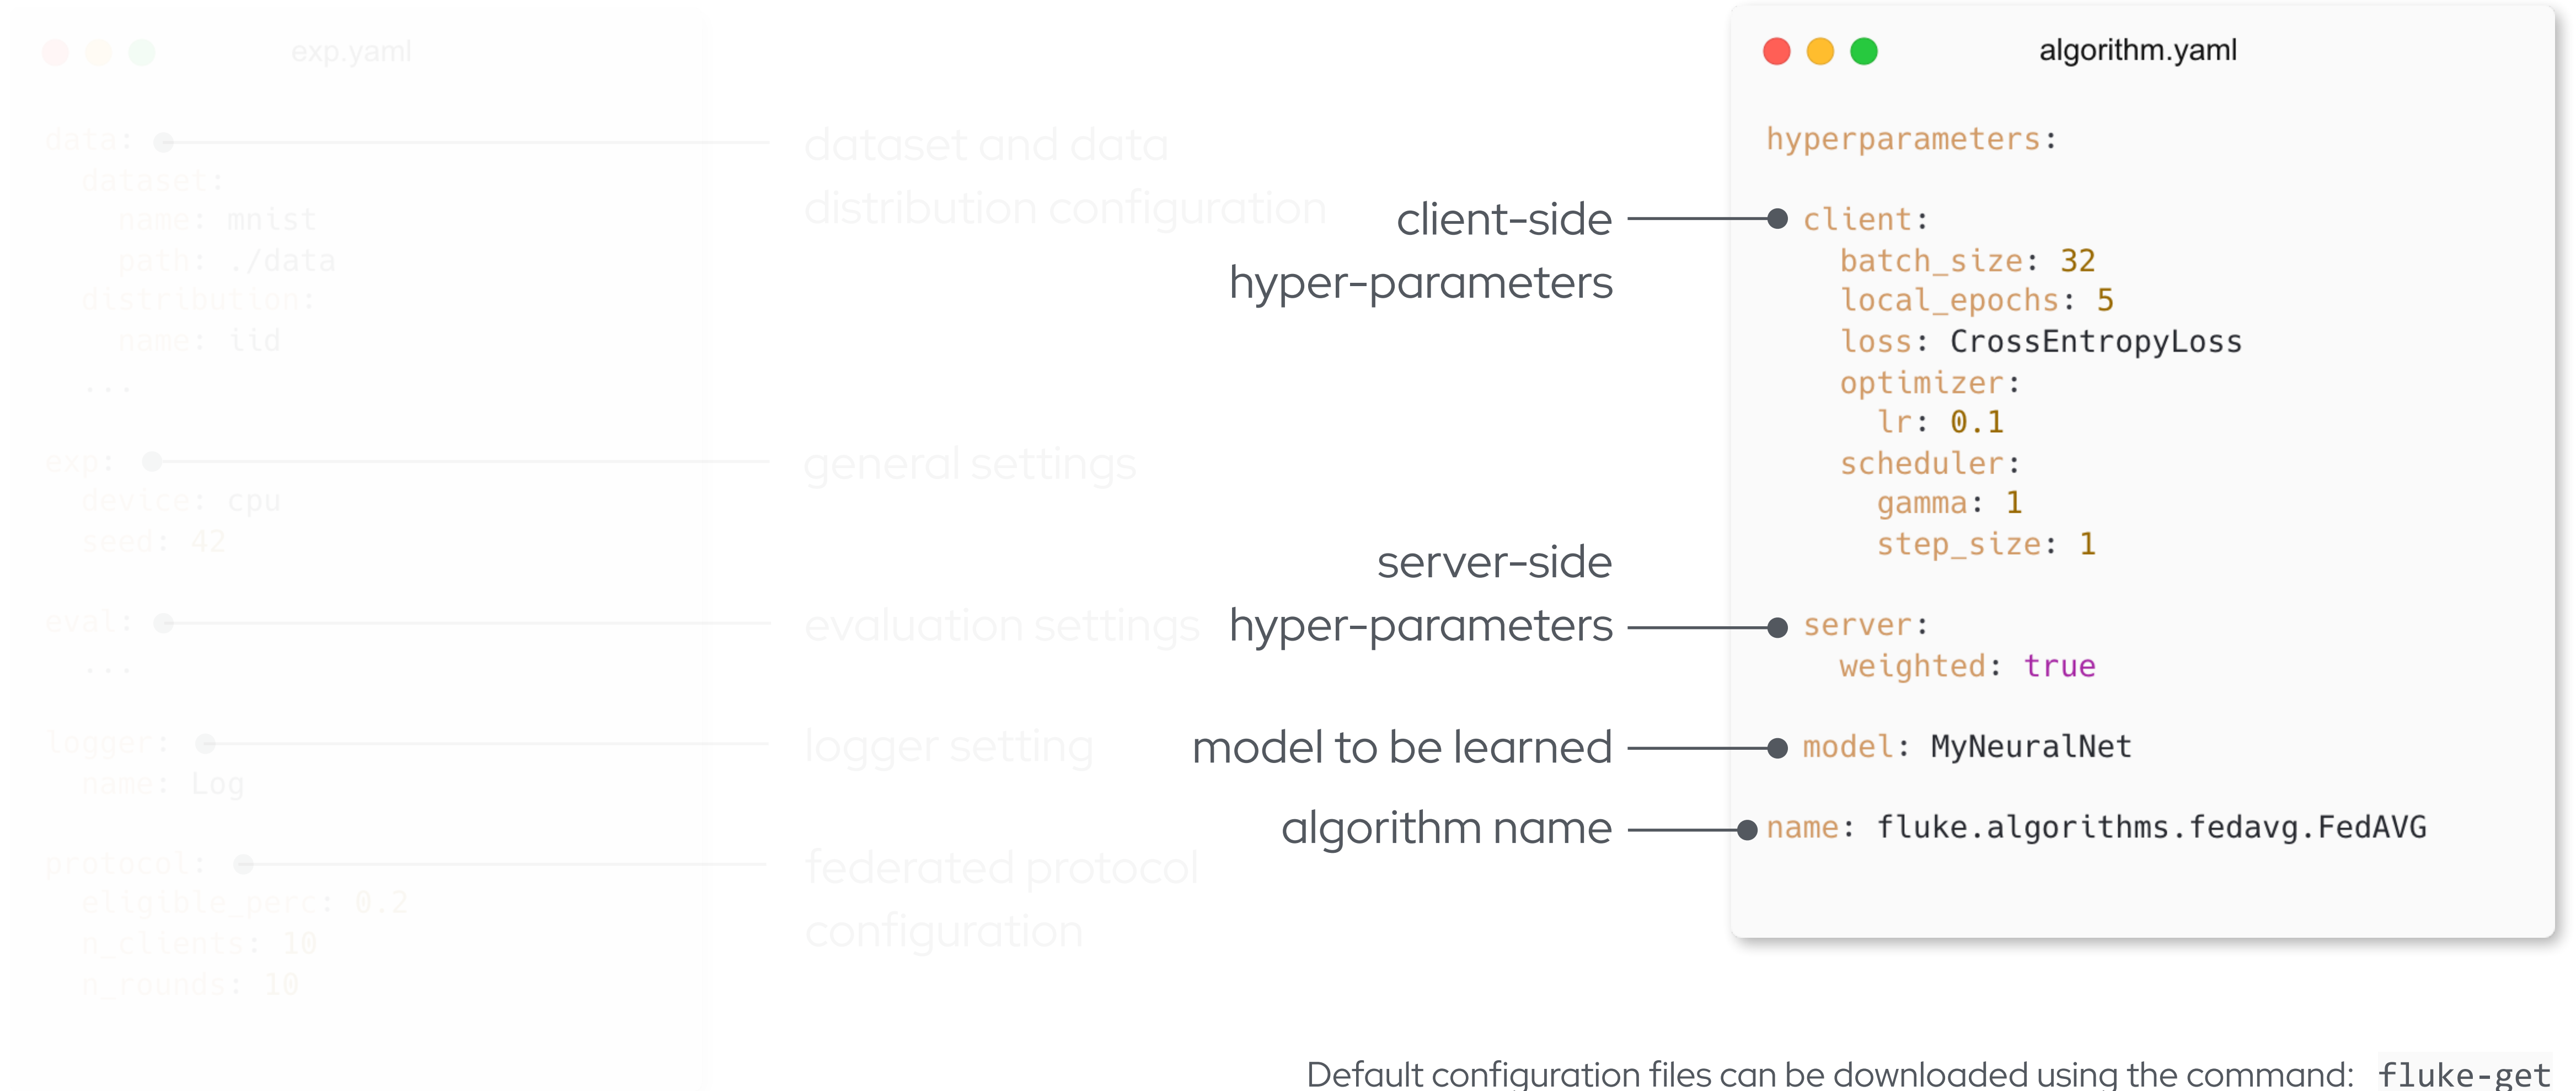
You can run your experiment outright with a single command!



(*) If you cloned the repo, the command (launched from the fluke folder) is `$ python -m fluke.run --config exp.yaml federation algorithm.yaml`

Configuration files

Two YAML files for everything you want to configure



Configuration files

Two YAML files for everything you want to configure

```
exp.yaml
data:
  dataset:
    name: mnist
    path: ./data
  distribution:
    name: iid
  ...
exp:
  device: cpu
  seed: 42
eval:
  ...
logger:
  name: Log
protocol:
  eligible_perc: 0.2
  n_clients: 10
  n_rounds: 10

algorithm.yaml
hyperparameters:
  client:
    batch_size: 32
    local_epochs: 5
    loss: CrossEntropyLoss
    optimizer:
      lr: 0.1
    scheduler:
      gamma: 1
      step_size: 1
  server:
    weighted: true
  model: MyNeuralNet
  name: fluke.algorithms.fedavg.FedAVG
```

dataset and data distribution configuration

general settings

evaluation settings

logger setting

federated protocol configuration

client-side hyper-parameters

server-side hyper-parameters

model to be learned

algorithm name

Default configuration files can be downloaded using the command: `fluke-get`

fluke CLI - not only federation

You can run the "same" experiment without the federation for comparison

same experiment but without the federation -
the number of epochs client-side are calculated*
as $n_rounds * eligible_perc * local_epochs$



```
$ fluke --config exp.yaml clients-only algorithm.yaml
```



```
$ fluke --config exp.yaml centralized algorithm.yaml
```

same type of experiment but with all the dataset centralised

(*) The number of epochs can be set by the user via an option of the command, e.g., `-epochs=100`

Example: FedAVG on MNIST

Fluke comes with many downloadable configuration files ready to be used/modified

downloads the default
experiment configuration file
(named `exp.yaml`) to `./config`

downloads the default fedavg
configuration file (named `fedavg.yaml`)
to `./config`

```

$ fluke-get config exp
$ fluke-get config fedavg
$ fluke --config ./config/exp.yaml federation ./config/fedavg.yaml


```

runs the federated algorithm
specified in `fedavg.yaml` on the
dataset specified in `exp.yaml`


(*) If you want to know all the available default configuration files use the command `$ fluke-get list`

f1uke logging

Performance can be logged on your preferred tool



communication cost*



classification
performance**

(e.g., accuracy, precision, recall, F1
- marco and micro)



system performance***




(*) The communication cost is estimated as the number of floating point numbers exchanged by the entire federation.

(**) Currently, fluke supports only classification but it is straightforward to extend to other tasks.

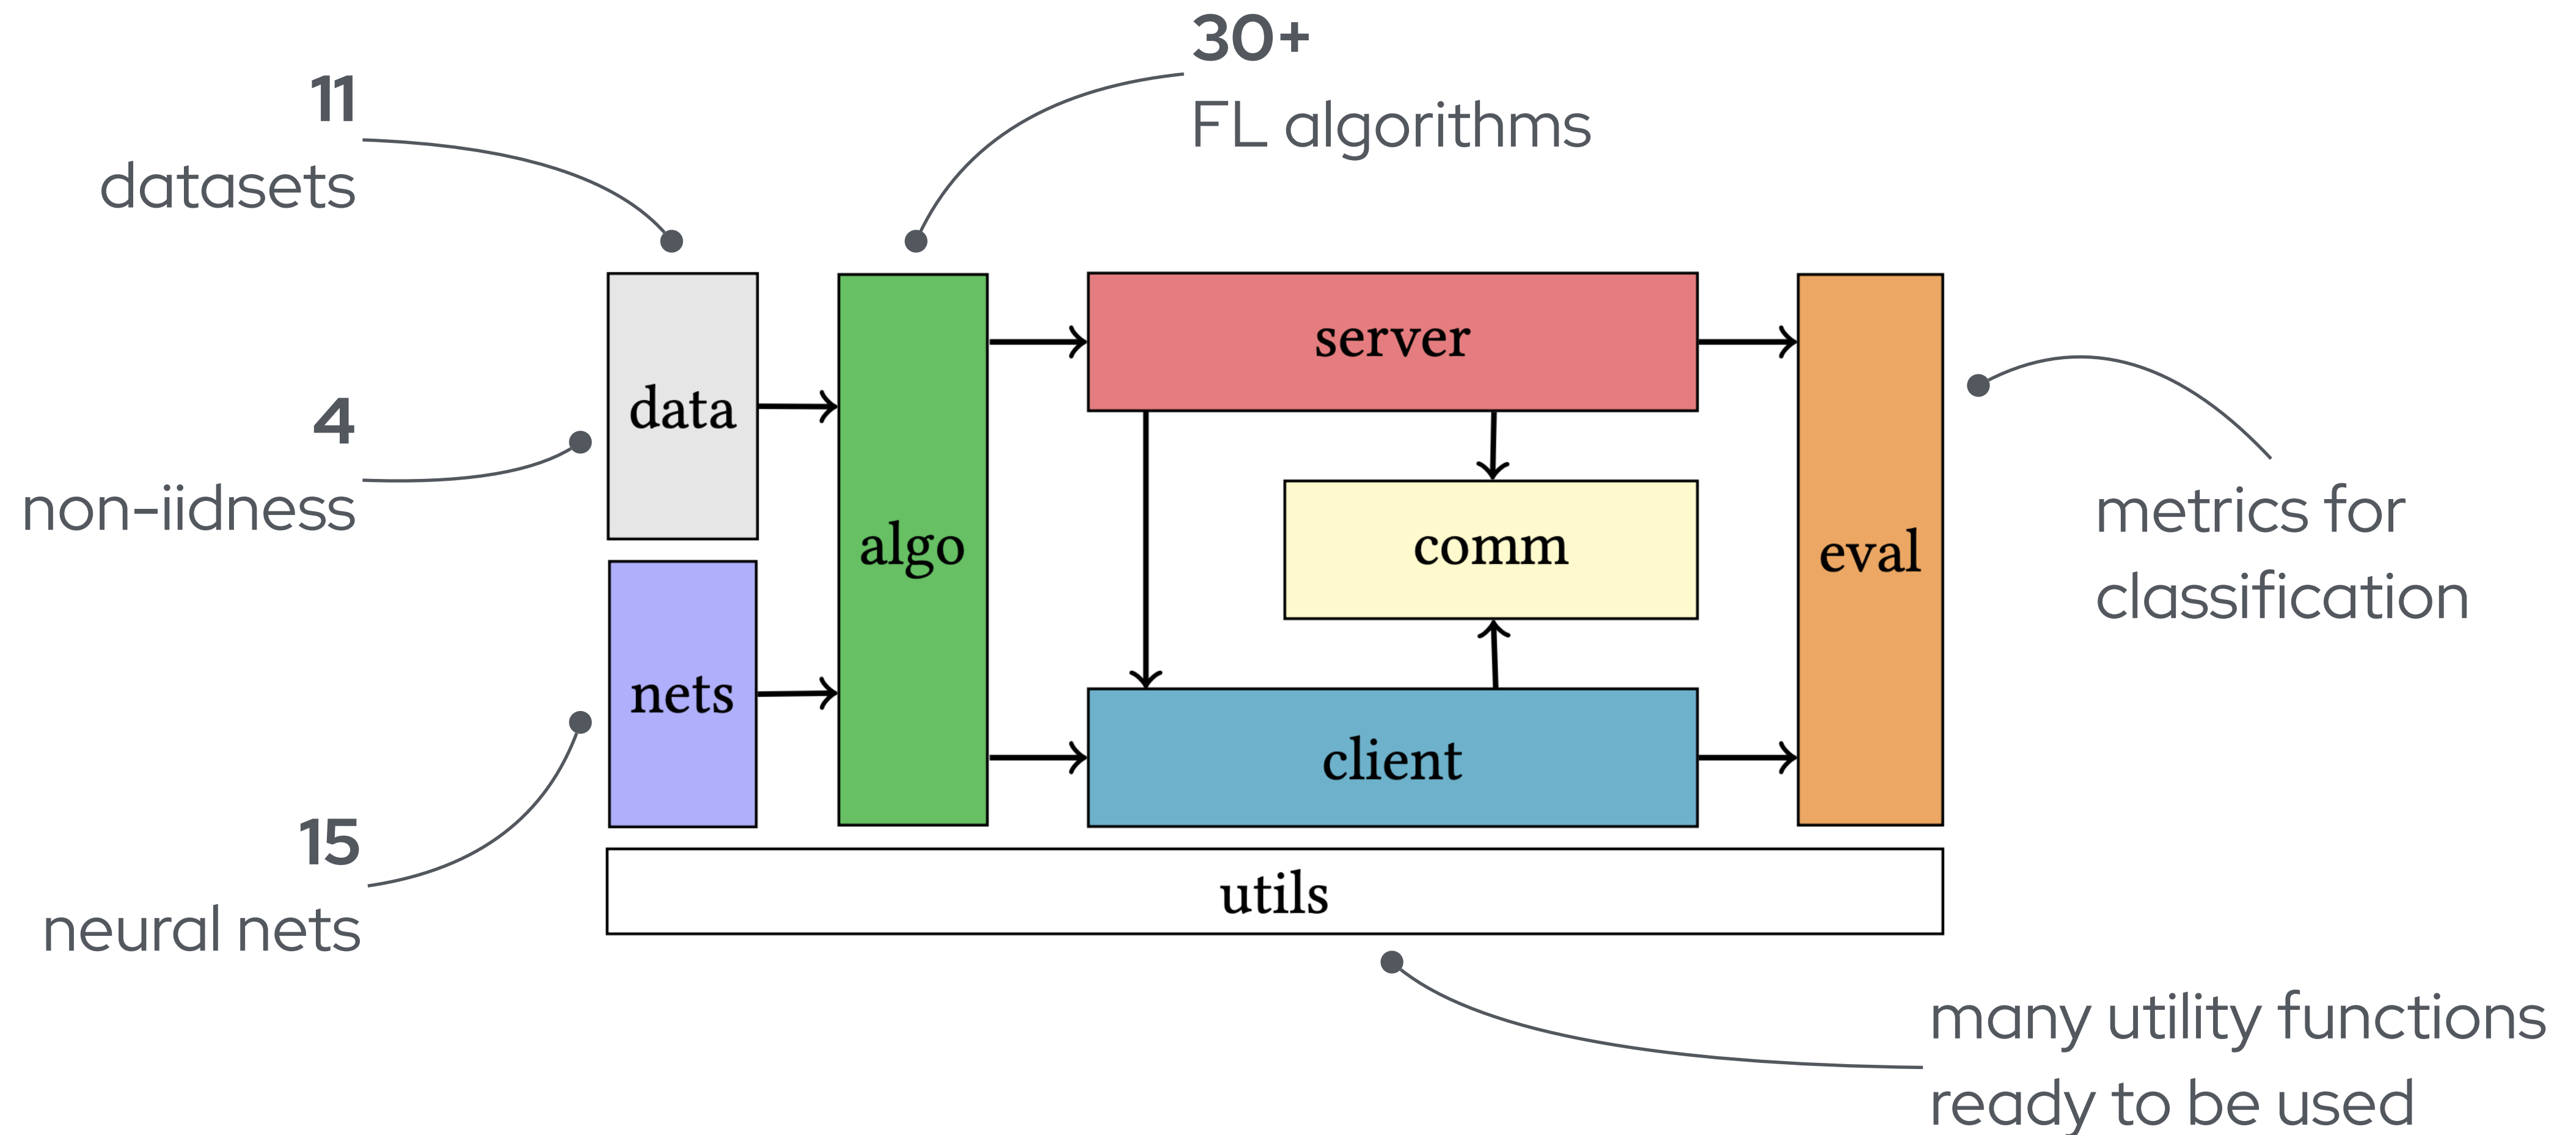
(***) System performance are automatically logged by W&B and ClearML.

f1uke python API

<https://github.com/makgyver/fluke>



The python API offers all you need to implement and test your FL ideas



f1uke server

Code readability is a key feature of f1uke

Algorithm 1 FederatedAveraging. The K clients are indexed by k ; B is the local minibatch size, E is the number of local epochs, and η is the learning rate.

Server executes:

initialize w_0

for each round $t = 1, 2, \dots$ do

$m \leftarrow \max(C \cdot K, 1)$

$S_t \leftarrow$ (random set of m clients)

for each client $k \in S_t$ in parallel do

$w_{t+1}^k \leftarrow \text{ClientUpdate}(k, w_t)$

$m_t \leftarrow \sum_{k \in S_t} n_k$

$w_{t+1} \leftarrow \sum_{k \in S_t} \frac{n_k}{m_t} w_{t+1}^k$

```
Server *
1 def fit(self, n_rounds: int, eligible_perc: float) -> None:
2
3     for round in range(n_rounds):
4         eligible = self.get_eligible_clients(eligible_perc)
5         self.broadcast_model(eligible)
6
7         for c, client in enumerate(eligible):
8             client.local_update(round + 1)
9
10        self.aggregate(eligible)
```

f1uke client

Code readability is a key feature of f1uke

ClientUpdate(k, w): // Run on client k
 $\mathcal{B} \leftarrow$ (split \mathcal{P}_k into batches of size B)
for each local epoch i from 1 to E **do**
 for batch $b \in \mathcal{B}$ **do**
 $w \leftarrow w - \eta \nabla \ell(w; b)$
return w to server

```
Client
1 def local_update(self, current_round: int):
2     self.receive_model()
3     self.fit()
4     self.send_model()
5
6 def fit(self, override_local_epochs: int = 0):
7     self.model.train()
8
9     if self.optimizer is None:
10         self.optimizer, self.scheduler = self.optimizer_cfg(self.model)
11
12     for _ in range(epochs):
13         for _, (X, y) in enumerate(self.train_set):
14             X, y = X.to(self.device), y.to(self.device)
15             self.optimizer.zero_grad()
16             y_hat = self.model(X)
17             loss = self.hyper_params.loss_fn(y_hat, y)
18             loss.backward()
19             self.optimizer.step()
20             self.scheduler.step()
21
22 def receive_model(self):
23     msg = self.channel.receive(self, self.server, msg_type="model")
24     self.model.load_state_dict(msg.payload.state_dict())
25
26 def send_model(self):
27     self.channel.send(Message(self.model, "model", self), self.server)
```

f1uke python API - Dataset loading & splitting



Loading the dataset

```
1 from fluke.data.datasets import Datasets
2 dataset = Datasets.get("mnist", path="./data")
```

dataset to (down)load*

folder where the dataset
will be stored/loaded



Splitting the dataset

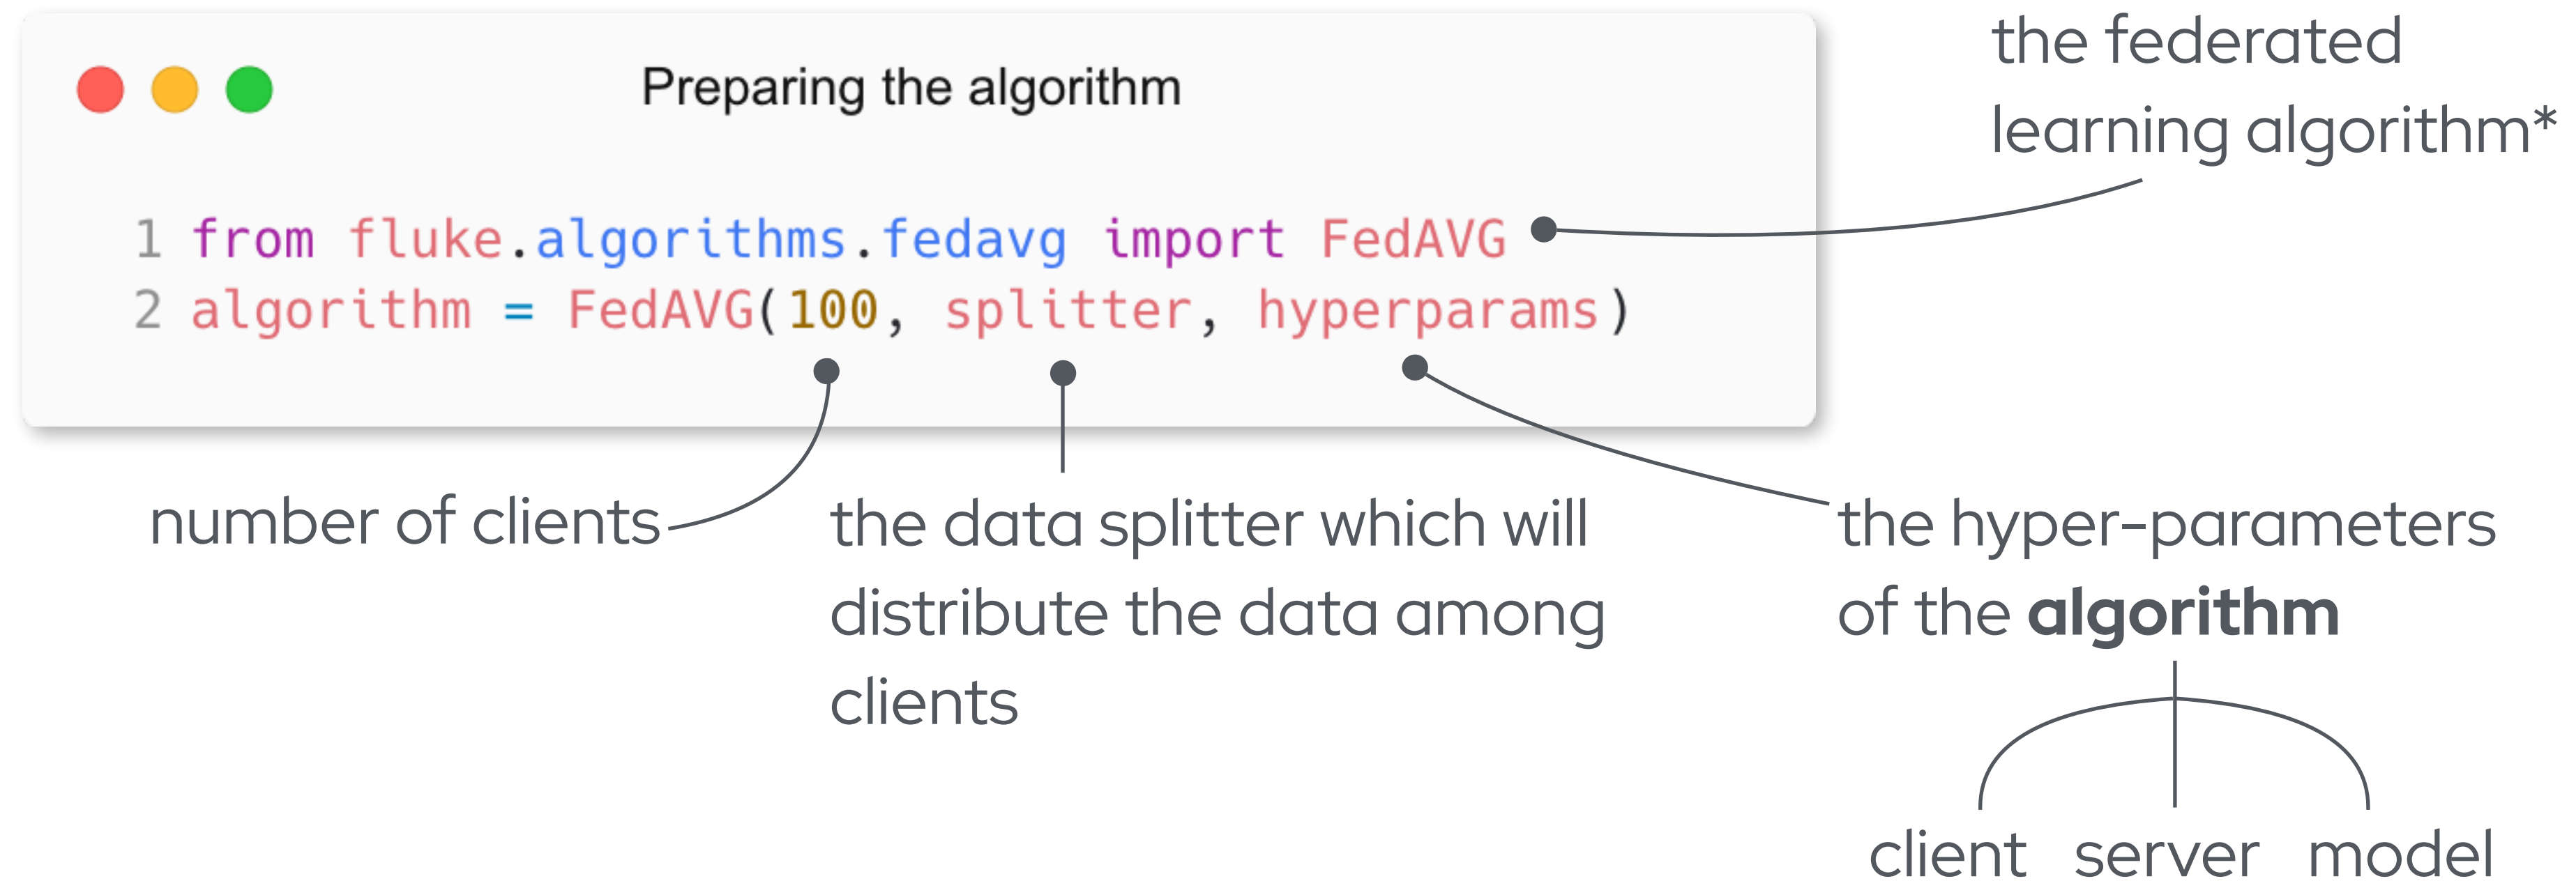
```
1 from fluke.data import DataSplitter
2 splitter = DataSplitter(dataset=dataset, distribution="iid")
```

the dataset to split

type of non-iidness

(*) fluke currently supports the following built-in datasets: MNIST, MNIST-M, SVHN, FEMNIST, EMNIST, CIFAR10, CIFAR100, Tiny Imagenet, Shakespeare, Fashion MNIST, and CINIC10.

f1uke python API - Federated algorithm



(*) fluke currently includes 31 different federated algorithms that you can use off-the-shelf.

f1uke python API - Hyper-parameters

You can load the hyper-parameters from file or hard-coding them

```
Hyper-parameters

1 from fluke import DDict
2
3 client_hp = DDict(
4     batch_size=10,
5     local_epochs=5,
6     loss="CrossEntropyLoss",
7     optimizer=DDict(name="SGD", lr=0.01)
8 )
9
10 server_hp = DDict(weighted=True)
11
12 # we put together the hyperparameters
13 hyperparams = DDict(client=client_hp,
14                     server=server_hp,
15                     model="MNIST_2NN")
```

```
Hyper-parameters from file

1 import yaml
2
3 with open("myconfig.yaml") as f:
4     config_alg = yaml.safe_load(f)
5
6 hyperparams = DDict(**config_alg)
```

client specific hyper-parameters,
including the optimizer and the
scheduler

server specific hyper-parameters

the shared model class*

(*) f1uke can also handle cases where clients and server own different models (like sub-network of the overall model)

fluke python API - Logging

Logging is handled using callbacks (i.e., design pattern Observer)

this is the default logger
(on console) but you can
also log on **W&B**,
Tensorboard or **ClearML**

the loggers in fluke only
log evaluation results

```
Logging
1 from fluke.utils.log import Log
2 logger = Log()
3 algorithm.set_callbacks(logger)
```

you can add all the observers you like.
you can observe the clients, the server
and the communication channel

f1uke python API - Start the training

```
Run the algorithm  
1 algorithm.run(n_rounds=2, eligible_perc=0.5)
```


Log on console

```
Round: 1  
{  
  'global': {  
    'accuracy': 0.8872,  
    'micro_precision': 0.8872,  
    'micro_recall': 0.8872,  
    'micro_f1': 0.8872,  
    'macro_precision': 0.88576,  
    'macro_recall': 0.88476,  
    'macro_f1': 0.88447  
  },  
  'comm_cost': 17811000  
}
```

Log on your preferred tool



You can save and restore the FL training



f1lukes - Adding a new FL algorithm

You just need to implement the core part of your algorithm as described in the paper

```
Define your client

1 from fluke.client import Client
2
3 class MyClient(Client):
4
5     def __init__(index: int,
6                 train_set: FastDataLoader,
7                 test_set: FastDataLoader,
8                 optimizer_cfg: OptimizerConfigurator,
9                 loss_fn: Module,
10                local_epochs: int = 3,
11                **kwargs: dict[str, Any]):
12         ...
13
14     def receive_model(self) -> None:
15         ...
16
17     def send_model(self) -> None:
18         ...
19
20     def fit(self, override_local_epochs: int = 0) -> float:
21         ...
22
```

```
Define your server

1 from fluke.server import Server
2
3 class Myserver(Server):
4
5     def __init__(self,
6                 model: torch.nn.Module,
7                 test_set: FastDataLoader,
8                 clients: Iterable[Client],
9                 weighted: bool = False,
10                **kwargs: dict[str, Any]):
11         ...
12
13     def aggregate(self, eligible: Iterable[Client]) -> None:
14         ...
```

```
Define your algorithm

1 class MyFLAlgo(CentralizedFL):
2
3     def get_server_class(self):
4         return MyServer
5
6     def get_client_class(self):
7         return MyClient
```

You just need to implement what characterise your FL algorithm

Implementing Kafè in fluke

Kafè has been presented at ECML PKDD 2024!

typical client selection process,
that is already implemented
in `fluke.server.Server*`

this must be implemented !!

this is the standard behaviour of a
FedAVG client, that is already
implemented in `fluke.client.Client`

Algorithm 1 KAFÈ

Require: T communication rounds, E local epochs, B local batchsize, h bandwidth of KDE, m number of clients participating in aggregation.

```
1: Server execute:
2:   Initialize model  $w^0$ 
3:   for  $t = 1, \dots, T$  do
4:      $m \leftarrow \max([C] \times K, 1)$ 
5:      $S_m \leftarrow$  random selection of  $m$  clients
6:     Send  $w^{(t-1)}$  to all clients.
7:     for chosen client  $k \in S_m$  in parallel do
8:        $w_k^{f,(t)}, w_k^{c,(t)} \leftarrow \text{LocalUpdating}(w^{(t-1)})$ 
9:
10:    Model aggregation:
11:       $w_g^{f,(t)} \leftarrow \sum_{k \in S_m} \frac{n_k}{n} w_k^{f,(t)}$ .
12:       $w_g^{c,(t)} \leftarrow \text{KAFÈ}(h, \frac{n_k}{n} w_k^{c,(t)})$ .
13:      Update  $w_g^{(t)} = (w_g^{f,(t)}, w_g^{c,(t)})$ 
14:    end for
15:  end for
16:
17:  LocalUpdating( $w^{(t-1)}$ ):
18:    for  $e = 1, 2, \dots, E - 1$  do
19:      for each batch  $B$  do
20:         $w_k^{(t)} \leftarrow w^{(t-1)} - \eta \nabla \ell(w^{(t-1)}, b)$ .
21:      end for
22:    end for
23:  return  $w_k^{(t)} = (w_k^{f,(t)}, w_k^{c,(t)})$ 
```

Implementing Kafè in fluke

Kafè is a FL algorithm presented at ECML PKDD 2024!

receive the client models $(w_k^{f,(t)}, w_k^{c,(t)})$

S_m

Model aggregation:

$$w_g^{f,(t)} \leftarrow \sum_{k \in S_m} \frac{n_k}{n} w_k^{f,(t)}$$

$$w_g^{c,(t)} \leftarrow \mathbf{KAFÈ}(h, \frac{n_k}{n} w_k^{c,(t)})$$

Update $w_g^{(t)} = (w_g^{f,(t)}, w_g^{c,(t)})$

```
Kafè Server

1 class KafeServer(Server):
  ~
  12
  13 def aggregate(self, eligible: Iterable[Client]) -> None:
  14     avg_model_sd = OrderedDict()
  15     clients_sd = self.get_client_models(eligible)
  16     weights = self._get_client_weights(eligible)
  17
  18     # get last layer of m clients' weights
  19     last_layer_weight_name = list(clients_sd[0].keys())[-2]
  20     last_layer_bias_name = list(clients_sd[0].keys())[-1]
  21
  22     for key in self.model.state_dict().keys():
  23         if key in (last_layer_weight_name, last_layer_bias_name):
  24             continue
  25         for i, client_sd in enumerate(clients_sd):
  26             if key not in avg_model_sd:
  27                 avg_model_sd[key] = weights[i] * client_sd[key]
  28             else:
  29                 avg_model_sd[key] = avg_model_sd[key] + weights[i] * client_sd[key]
  30
```

Implementing Kafè in fluke

Kafè is a FL algorithm presented at ECML PKDD 2024!

$$w_g^{c,(t)} \leftarrow \mathbf{KAFÈ}(h, \frac{n_k}{n} w_k^{c,(t)}).$$

Classification layers aggregation. The m classification layers of the m local models can be denoted as $\{w_1^{c,(t)}, w_2^{c,(t)}, \dots, w_k^{c,(t)}, \dots\}$. To evaluate the probability density function $\hat{f}(\cdot)$ for these classification layers w^c , we employ KDE as follows:

$$\hat{f}(w^c) = \frac{1}{mh^d} \sum_{k \in S_m^{(t)}} K\left(\frac{w_k^{c,(t)} - w^c}{h}\right) \quad (5)$$

where d denotes the dimension of the samples, similar to Eq. (1). In this context, d corresponds to the dimension of $w_k^{c,(t)}$ after flattening. To account for the varying contributions of the classification layers from different local models, we introduce weighting factors $\frac{n_k}{n}$ to the KDE formulation. Consequently, Eq. (5) is transformed into Eq. (6).

$$\hat{f}(w^c) = \frac{1}{h^d} \sum_{k \in S_m^{(t)}} \frac{n_k}{n} K\left(\frac{w_k^{c,(t)} - w^c}{h}\right) \quad (6)$$

Next, by sampling m new samples from the KDE and averaging them, we obtain the new classification layer $w_g^{c,(t)}$ for the global model. Lastly, the global

```
Kafè Server

1 class KafèServer(Server):
2
3     def __init__(self,
4                 model: torch.nn.Module,
5                 test_set: FastDataLoader,
6                 clients: Iterable[Client],
7                 weighted: bool = False,
8                 bandwidth: float = 1.0):
9         super().__init__(model=model, test_set=test_set,
10                          clients=clients, weighted=weighted)
11         self.hyper_params.update(bandwidth=bandwidth)
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31         w_last_layer = []
32         b_last_layer = []
33
34         for csd in clients_sd:
35             w_last_layer.append(np.array(csd[last_layer_weight_name]))
36             b_last_layer.append(np.array(csd[last_layer_bias_name]))
37
38         w_last_layer = np.array(w_last_layer).reshape(len(w_last_layer), -1)
39         b_last_layer = np.array(b_last_layer).reshape(len(b_last_layer), -1)
40
41         # using KDE get the kernel density of last layers
42         kde_w = KernelDensity(kernel='gaussian',
43                               bandwidth=self.hyper_params.bandwidth).fit(w_last_layer,
44                                                                              sample_weight=weights)
45         kde_b = KernelDensity(kernel='gaussian',
46                               bandwidth=self.hyper_params.bandwidth).fit(b_last_layer,
47                                                                              sample_weight=weights)
48
49         # sample m samples and average, then obtain a new last layer for the global model
50         w_last_layer_new = np.mean(kde_w.sample(len(w_last_layer)), axis=0)
51         b_last_layer_new = np.mean(kde_b.sample(len(b_last_layer)), axis=0)
52
53         # update last layer
54         avg_model_sd[last_layer_weight_name] = torch.tensor(w_last_layer_new.reshape(
55             clients_sd[0][last_layer_weight_name].shape))
56         avg_model_sd[last_layer_bias_name] = torch.tensor(b_last_layer_new.reshape(
57             clients_sd[0][last_layer_bias_name].shape))
58
59         self.model.load_state_dict(avg_model_sd)
```

Implementing Kafè in fluke

Kafè is a FL algorithm presented at ECML PKDD 2024!

```
Kafè Server

1 class Kafeserver(Server):
2
3     def __init__(self,
4                 model: torch.nn.Module,
5                 test_set: FastDataLoader,
6                 clients: Iterable[Client],
7                 weighted: bool = False,
8                 bandwidth: float = 1.0):
9         super().__init__(model=model, test_set=test_set,
10                          clients=clients, weighted=weighted)
11         self.hyper_params.update(bandwidth=bandwidth)
```

$$w_g^{c,(t)} \leftarrow \mathbf{KAFÈ}(h, \frac{n_k}{n} w_k^{c,(t)}).$$


Implementing Kafè in fluke

Kafè is a FL algorithm presented at ECML PKDD 2024!

$$w_g^{c,(t)} \leftarrow \mathbf{KAFÈ}(h, \frac{n_k}{n} w_k^{c,(t)}).$$

Classification layers aggregation. The m classification layers of the m local models can be denoted as $\{w_1^{c,(t)}, w_2^{c,(t)}, \dots, w_k^{c,(t)}, \dots\}$. To evaluate the probability density function $\hat{f}(\cdot)$ for these classification layers w^c , we employ KDE as follows:

$$\hat{f}(w^c) = \frac{1}{mh^d} \sum_{k \in S_m^{(t)}} K\left(\frac{w_k^{c,(t)} - w^c}{h}\right) \quad (5)$$

where d denotes the dimension of the samples, similar to Eq. (1). In this context, d corresponds to the dimension of $w_k^{c,(t)}$ after flattening. To account for the varying contributions of the classification layers from different local models, we introduce weighting factors $\frac{n_k}{n}$ to the KDE formulation. Consequently, Eq. (5) is transformed into Eq. (6).

$$\hat{f}(w^c) = \frac{1}{h^d} \sum_{k \in S_m^{(t)}} \frac{n_k}{n} K\left(\frac{w_k^{c,(t)} - w^c}{h}\right) \quad (6)$$

Next, by sampling m new samples from the KDE and averaging them, we obtain the new classification layer $w_g^{c,(t)}$ for the global model. Lastly, the global

```
Kafè Server

1 class KafèServer(Server):
2
3     def __init__(self,
4                 model: torch.nn.Module,
5                 test_set: FastDataLoader,
6                 clients: Iterable[Client],
7                 weighted: bool = False,
8                 bandwidth: float = 1.0):
9         super().__init__(model=model, test_set=test_set,
10                          clients=clients, weighted=weighted)
11         self.hyper_params.update(bandwidth=bandwidth)
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31         w_last_layer = []
32         b_last_layer = []
33
34         for csd in clients_sd:
35             w_last_layer.append(np.array(csd[last_layer_weight_name]))
36             b_last_layer.append(np.array(csd[last_layer_bias_name]))
37
38         w_last_layer = np.array(w_last_layer).reshape(len(w_last_layer), -1)
39         b_last_layer = np.array(b_last_layer).reshape(len(b_last_layer), -1)
40
41         # using KDE get the kernel density of last layers
42         kde_w = KernelDensity(kernel='gaussian',
43                               bandwidth=self.hyper_params.bandwidth).fit(w_last_layer,
44                                                                              sample_weight=weights)
45         kde_b = KernelDensity(kernel='gaussian',
46                               bandwidth=self.hyper_params.bandwidth).fit(b_last_layer,
47                                                                              sample_weight=weights)
48
49         # sample m samples and average, then obtain a new last layer for the global model
50         w_last_layer_new = np.mean(kde_w.sample(len(w_last_layer)), axis=0)
51         b_last_layer_new = np.mean(kde_b.sample(len(b_last_layer)), axis=0)
52
53         # update last layer
54         avg_model_sd[last_layer_weight_name] = torch.tensor(w_last_layer_new.reshape(
55             clients_sd[0][last_layer_weight_name].shape))
56         avg_model_sd[last_layer_bias_name] = torch.tensor(b_last_layer_new.reshape(
57             clients_sd[0][last_layer_bias_name].shape))
58
59         self.model.load_state_dict(avg_model_sd)
```


Implementing Kafè in fluke

Kafè is a FL algorithm presented at ECML PKDD 2024!

$$w_g^{c,(t)} \leftarrow \mathbf{KAFÈ}(h, \frac{n_k}{n} w_k^{c,(t)}).$$

Classification layers aggregation. The m classification layers of the m local models can be denoted as $\{w_1^{c,(t)}, w_2^{c,(t)}, \dots, w_k^{c,(t)}, \dots\}$. To evaluate the probability density function $\hat{f}(\cdot)$ for these classification layers w^c , we employ KDE as follows:

$$\hat{f}(w^c) = \frac{1}{mh^d} \sum_{k \in S_m^{(t)}} K\left(\frac{w_k^{c,(t)} - w^c}{h}\right) \quad (5)$$

where d denotes the dimension of the samples, similar to Eq. (1). In this context, d corresponds to the dimension of $w_k^{c,(t)}$ after flattening. To account for the varying contributions of the classification layers from different local models, we introduce weighting factors $\frac{n_k}{n}$ to the KDE formulation. Consequently, Eq. (5) is transformed into Eq. (6).

$$\hat{f}(w^c) = \frac{1}{h^d} \sum_{k \in S_m^{(t)}} \frac{n_k}{n} K\left(\frac{w_k^{c,(t)} - w^c}{h}\right) \quad (6)$$

Next, by sampling m new samples from the KDE and averaging them, we obtain the new classification layer $w_g^{c,(t)}$ for the global model. Lastly, the global

```
Kafè Server

1 class KafèServer(Server):
2
3     def __init__(self,
4                 model: torch.nn.Module,
5                 test_set: FastDataLoader,
6                 clients: Iterable[Client],
7                 weighted: bool = False,
8                 bandwidth: float = 1.0):
9         super().__init__(model=model, test_set=test_set,
10                          clients=clients, weighted=weighted)
11         self.hyper_params.update(bandwidth=bandwidth)
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31         w_last_layer = []
32         b_last_layer = []
33
34         for csd in clients_sd:
35             w_last_layer.append(np.array(csd[last_layer_weight_name]))
36             b_last_layer.append(np.array(csd[last_layer_bias_name]))
37
38         w_last_layer = np.array(w_last_layer).reshape(len(w_last_layer), -1)
39         b_last_layer = np.array(b_last_layer).reshape(len(b_last_layer), -1)
40
41         # using KDE get the kernel density of last layers
42         kde_w = KernelDensity(kernel='gaussian',
43                               bandwidth=self.hyper_params.bandwidth).fit(w_last_layer,
44                                                                              sample_weight=weights)
45         kde_b = KernelDensity(kernel='gaussian',
46                               bandwidth=self.hyper_params.bandwidth).fit(b_last_layer,
47                                                                              sample_weight=weights)
48
49         # sample m samples and average, then obtain a new last layer for the global model
50         w_last_layer_new = np.mean(kde_w.sample(len(w_last_layer)), axis=0)
51         b_last_layer_new = np.mean(kde_b.sample(len(b_last_layer)), axis=0)
52
53         # update last layer
54         avg_model_sd[last_layer_weight_name] = torch.tensor(w_last_layer_new.reshape(
55             clients_sd[0][last_layer_weight_name].shape))
56         avg_model_sd[last_layer_bias_name] = torch.tensor(b_last_layer_new.reshape(
57             clients_sd[0][last_layer_bias_name].shape))
58
59         self.model.load_state_dict(avg_model_sd)
```

Implementing Kafè in fluke

Kafè is a FL algorithm presented at ECML PKDD 2024!

Next, by sampling m new samples from the KDE and averaging them, we obtain the new classification layer $w_g^{c,(t)}$ for the global model. Lastly, the global

```
w_last_layer_new = np.mean(kde_w.sample(len(w_last_layer)), axis=0)
b_last_layer_new = np.mean(kde_b.sample(len(b_last_layer)), axis=0)

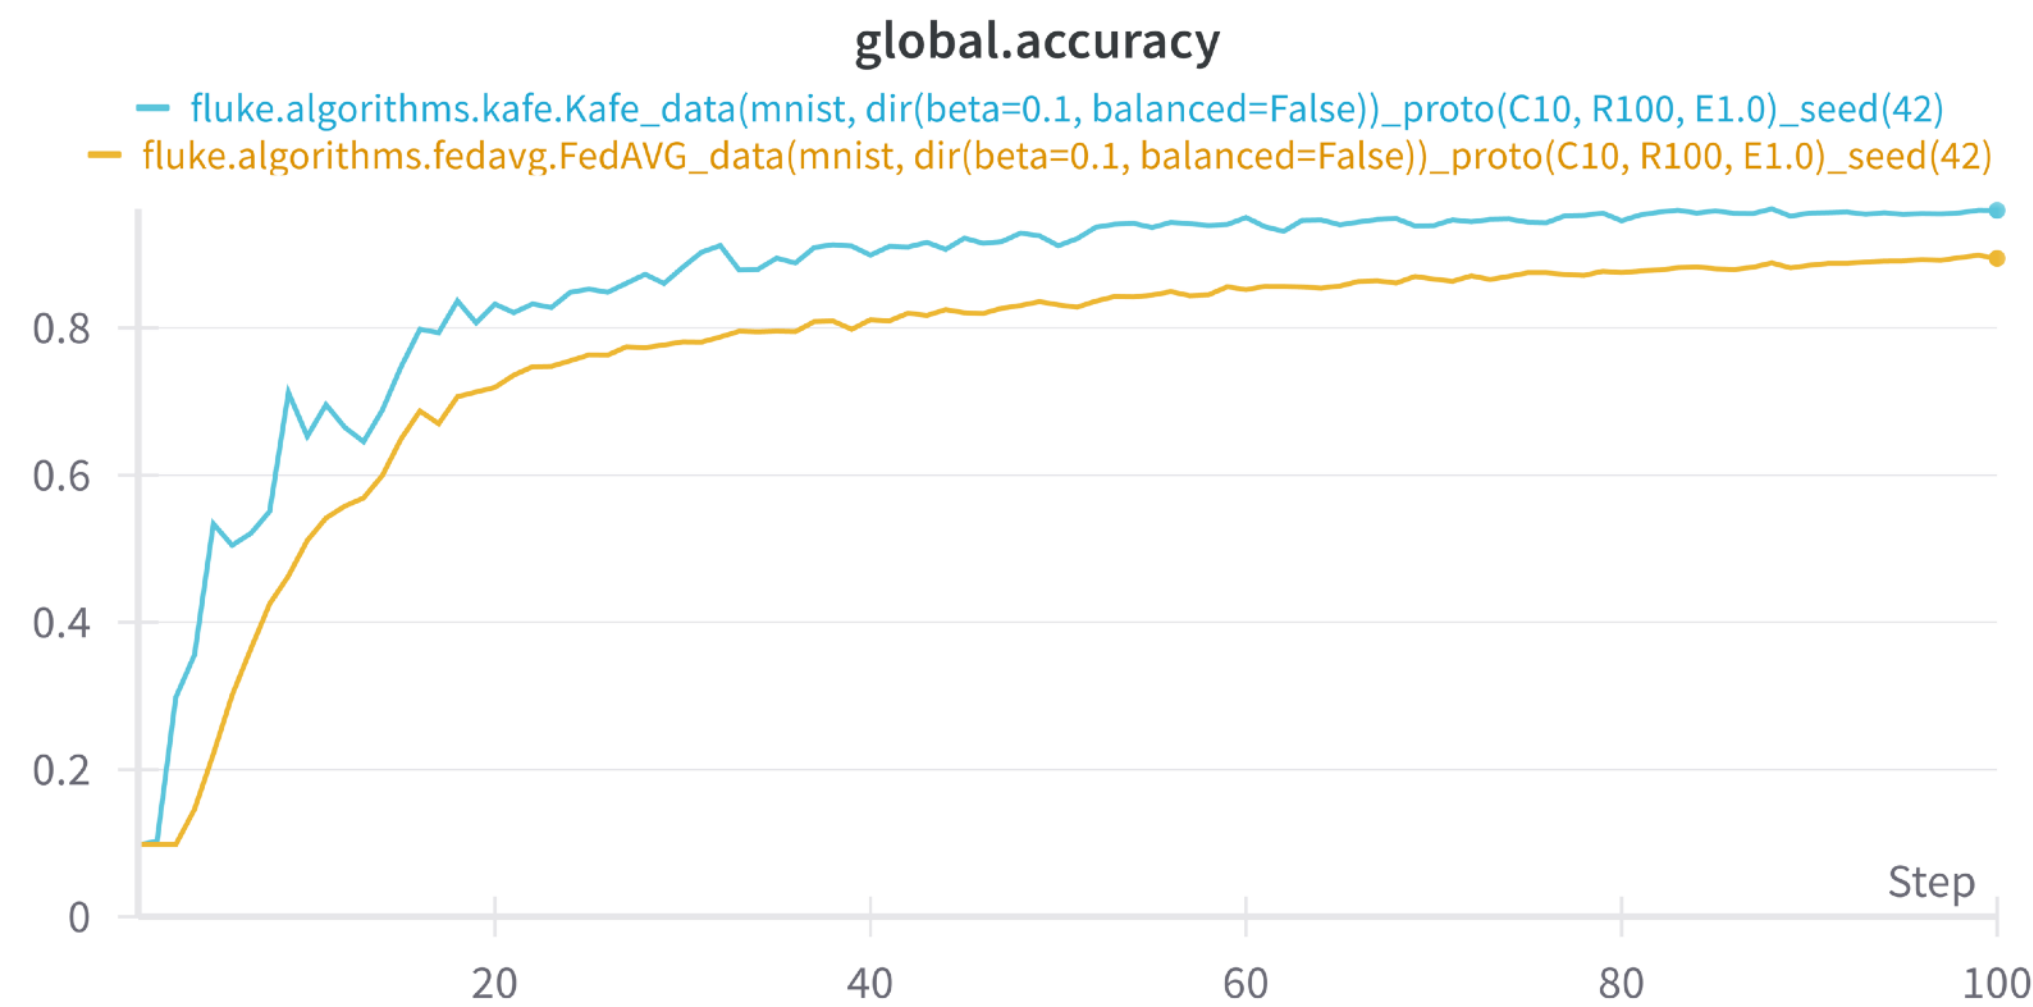
# update last layer
avg_model_sd[last_layer_weight_name] = torch.tensor(w_last_layer_new.reshape(
    clients_sd[0][last_layer_weight_name].shape))
avg_model_sd[last_layer_bias_name] = torch.tensor(b_last_layer_new.reshape(
    clients_sd[0][last_layer_bias_name].shape))

self.model.load_state_dict(avg_model_sd)
```


Testing Kafè in fluke

You just need to define the configuration files and use the fluke CLI

```
Running Kafè  
1 $ fluke --config exp.yaml federation kafe.yaml
```



fluke results (W&B plot)



Paper's result

Give it a try, you won't regret it!

fluke 0.3.0 is now available!

30+

FL algorithms

11

datasets

15

neural networks



propose your FL algorithm
to be included in fluke!



<https://github.com/makgyver/fluke>

<https://makgyver.github.io/fluke/>

Credits

All the icons have been downloaded from www.flaticon.com

[quick] Icon made by Cuputo from www.flaticon.com

[chart] Icon made by Freepik from www.flaticon.com

[idea] Icon made by Freepik from www.flaticon.com

[network cost] Icon made by juicy_fish from www.flaticon.com

[performance] Icon made by juicy_fish from www.flaticon.com

[programming] Icon made by juicy_fish from www.flaticon.com

[close] Icon made by Pixel Perfect from www.flaticon.com

[check-mark] Icon made by Ian June from www.flaticon.com

[pause] Icon made by bqlqn from www.flaticon.com

[jigsaw] Icon made by monkik from www.flaticon.com