
Federated Hypergradient Descent

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 In this work, we explore combining automatic hyperparameter tuning and opti-
2 mization for federated learning (FL) in an online, one-shot procedure. We apply
3 a principled approach on a method for adaptive client learning rate, number of
4 local steps, and batch size. In our federated learning applications, our primary
5 motivations are minimizing communication budget as well as local computational
6 resources in the training pipeline. Conventionally, hyperparameter tuning meth-
7 ods involve at least some degree of trial-and-error, which is known to be sample
8 inefficient. In order to address our motivations, we propose FATHOM (Federated
9 AuTomatic Hyperparameter OptiMization) as a one-shot online procedure. We
10 investigate the challenges and solutions of deriving analytical gradients with respect
11 to the hyperparameters of interest. Our approach is inspired by the fact that we
12 have full knowledge of all components involved in our training process, and this
13 fact can be exploited in our algorithm impactfully. We show that FATHOM is
14 more communication efficient than Federated Averaging (FedAvg) with optimized,
15 static valued hyperparameters, and is also more computationally efficient overall.
16 As a communication efficient, one-shot online procedure, FATHOM solves the
17 bottleneck of costly communication and limited local computation, by eliminat-
18 ing a potentially wasteful tuning process, and by optimizing the hyperparamters
19 adaptively throughout the training procedure without trial-and-error. We show
20 our numerical results through extensive empirical experiments with the Federated
21 EMNIST-62 (FEMNIST) and Federated Stack Overflow (FSO) datasets, using
22 FedJAX as our baseline framework.

23 1 Introduction

24 Federated learning (FL) for on-device applications has its obvious social implications, due to its
25 inherent privacy-protection feature. It opens up a broad range of opportunities to allow a massive
26 number of devices to collaborate in developing a shared model by retaining private data on the
27 devices. The ubiquity of machine learning (ML) on consumer data, coupled with the growth of privacy
28 concerns, has pushed researchers and developers to look for new ways to protect and benefit end-users.
29 In order for FL to deliver its promise in deployed applications, there are still many open challenges
30 remained to be solved. We are especially interested in the overall communication efficiency of the FL
31 pipeline for it to be realistically deployed in a unique communication environment over expensive
32 links. To begin, consider a typical step in a machine learning (ML) pipeline: hyperparameter tuning.
33 Whether it is in a centralized, distributed or federated setting, it is an essential step to achieve an
34 optimal operation for the training process. At the heart of an ML training process is the optimization
35 algorithm. In particular, we are interested in using Federated Averaging (FedAvg) as our baseline

36 federated optimization algorithm for our work. This is because, despite all the recent innovations in
37 FL since its introduction in 2016 by [McMahan et al. \[2016\]](#), FedAvg remains the de facto standard in
38 federated optimization for both research and practice, due to its simplicity and empirical effectiveness.
39 In order for FedAvg to operate effectively, it requires properly tuned hyperparameter values.

40 Our work focuses specifically on hyperparameter optimization (HPO) of: 1) client learning rate,
41 2) number of local steps, as well as 3) batch size, for FedAvg. We propose FATHOM (Federated
42 AuTomatic Hyperparameter OptiMization), which is an online algorithm that operates as a one-shot
43 procedure. In the rest of this paper, we will go through a few notable recent state-of-the-art works on
44 this topic, and make justifications for our new approach. Then we will derive a few key steps for our
45 algorithm, followed by a theoretical convergence bound for adaptive learning rate and number of local
46 steps in the non-convex regime. Lastly, we present numerical results on our empirical experiments
47 with neural networks on the FEMNIST and FSO datasets.

48 Our contributions are as follows:

- 49 • We derive gradients with respect to client learning rate and number of local steps for FedAvg,
50 for an online optimization procedure. We propose FATHOM, a practical one-shot procedure
51 for joint-optimization of hyperparameters and model parameters, for FedAvg.
- 52 • We derive a new convergence upper-bound with a relaxed condition (see Section 4 and
53 remark 2), to highlight the benefits from the extra degree-of-freedom that FATHOM delivers
54 for performance gains.
- 55 • We present empirical results that show state-of-the-art performance. To our knowledge,
56 we are the first to show gain from an online HPO procedure over a well-tuned equivalent
57 procedure with fixed hyperparameter values.

58 2 Related Work and Justifications for FATHOM

59 We explore the question whether the FATHOM approach is justified over the more recent, state-of-
60 the-art methods that are designed for the same goal: a single-shot online hyperparameter optimization
61 procedure for FL. [Zhou et al. \[2022\]](#) proposed Federated Loss Surface Aggregation (FLoRA), a
62 general single-shot HPO for FL, which works by treating HPO as a black-box problem and by
63 performing loss surface aggregation for training the global model. [Khodak et al. \[2021\]](#) draws
64 inspiration from weight-sharing in Neural Architectural Search ([Pham et al. \[2018\]](#), [Cai et al. \[2019\]](#)),
65 and proposed FedEx, which is an online hyperparameter tuning algorithm that uses exponentiated
66 gradients to update hyperparameters. On the other hand, [Mostafa \[2019\]](#)'s RMAH and [Guo et al.
67 \[2022\]](#)'s Auto-FedRL both use REINFORCE ([Williams \[1992\]](#)) in their agents to update hyperparam-
68 eters in an online manner, by using relative loss as their trial rewards. One basic assumption among
69 these methods, is that at least some of the gradients with respect to the hyperparameters are unavail-
70 able directly. Generalized techniques are used to update these quantities, involving Monte-Carlo
71 sampling and evaluation with held-out data. One key benefit with techniques such as these is their
72 generalizability for a wide range of different hyperparameters. On the other hand, we identify a few
73 areas with these methods that we would like to improve on. One, information about the internals of
74 the procedure can and should be exploited. Two, communication overhead becomes a concern, since
75 sufficient Monte-Carlo sampling is required for some of these techniques to converge, an example
76 being the re-parametrization trick ([Kingma and Welling \[2013\]](#)) which is used for FedEx, RMAH and
77 Auto-FedRL. From initial observations of their empirical results, while these methods are successful
78 in hyperparameter tuning and reaching target model accuracy as shown in these works, these goals are
79 achieved in unspecified numbers of total communication rounds from works based on RL approaches
80 such as [Mostafa \[2019\]](#) and [Guo et al. \[2022\]](#).

81 The above observations justify exploring our problem differently from previous approaches. Our
82 method exploits full knowledge of the training process, and it does not require sufficient trials
83 at potential expense of communication budget. Inspired by the hypergradient descent techniques
84 developed by [Baydin et al. \[2017\]](#) and [Amid et al. \[2022\]](#) for centralized optimization learning rate,
85 we develop FATHOM by directing deriving analytical gradients with respect to the hyperparameters
86 of interest. The result is a sample efficient method which offers both improvements in communication
87 efficiency and reduced local computation in a single-shot online optimization procedure. Meanwhile,
88 FATHOM is not as flexibly applicable in optimizing a wide range of hyperparameters, since each

89 gradient needs to be derived separately to take advantage of our full knowledge of the training process.
 90 We believe this approach is a performance advantage, at the expense of its flexibility.

91 There are other notable relevant works. Charles and Konečný [2020] and Li et al. [2019] proved that
 92 reducing the client learning rate during training is necessary to reach the true objective. Yet, a line of
 93 interesting works, such as Dai et al. [2020] and Holly et al. [2021] applies Bayesian Optimization
 94 (BO) on federated hyperparameter tuning, by treating it as a closed-box optimization problem. Dai
 95 et al. [2021] further updates their use of BO in FL by incorporating differential privacy. However,
 96 these BO-based works do not consider adaptive hyperparameters. Yet, another work (Wang and Joshi
 97 [2018]) shares similarity to our approach of optimally adapting the number of local steps, with their
 98 adaptive communication strategy, AdaComm, in the distributed setting. However, their main interest
 99 is reducing wall-clock time. Lastly, around the same time of this writing, Wang et al. [2022] publishes
 100 their benchmark suite for FL HPO, called FedHPO-B, which would be valuable to our future work.

101 3 Methodology

102 In this section we formalize the problem of hyperparameter optimization (HPO) for FL. We first
 103 review FedAvg, a de facto standard of federated optimization methods for research baseline and
 104 practice. Then, we present our method for online-tuning of its hyperparameters, specifically client
 105 learning rate, number of local steps, and batch size. We call our method FATHOM (Federated
 106 AuTomatic Hyperparameter OptiMization).

107 3.1 Problem Definition

108 In this paper, we consider the empirical risk minimization (ERM) across all the client data, as an
 109 unconstrained optimization problem:

$$f^* := \min_{x \in \mathbb{R}^d} \left[f(x) := \frac{1}{m} \sum_{i=1}^m f_i(x) \right] \quad (1)$$

110 where $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is the loss function for data stored in local client index i with d being the
 111 dimension of the parameters x , m is number of clients, and $f^* = f(x_*)$ where x_* is a stationary
 112 solution to the ERM problem in eq(1).

113 To facilitate some of the discussions that follow, it helps to define assumptions here as we do
 114 throughout the rest of this paper:

115 **Assumption 1.** (*Unbiased Local Gradient Estimator*) Let $g_i(x)$ be the unbiased, local gradient
 116 estimator of $\nabla f_i(x)$, i.e., $\mathbb{E}[g_i(x)] = \nabla f_i(x)$, $\forall x$, and $i \in [m]$.

117 3.2 Federated Optimization and Tuning of Hyperparameters

118 **Federated Averaging (FedAvg)** We describe the operations of FedAvg from McMahan et al.
 119 [2016], as follows. At any round t , each of the m clients takes a total of K_i local SGD steps, where
 120 $K_i = \lfloor E\nu_i/B \rfloor$, and where ν_i is the number of data samples from client index i , B is batch size,
 121 with epoch number $E = 1$ being a common baseline. In this version of FedAvg, heterogeneous data
 122 size is accommodated across clients, and the number of local steps can be manipulated via E and
 123 B as hyperparameters. Each local SGD step updates the local model parameters of each client i as
 124 follows: $x_{t,k+1}^i = x_{t,k}^i - \eta_L g_i(x_{t,k}^i)$, where η_L is the local learning rate and $k \in [K]$ is the local step
 125 index. To conclude each round, these clients return the local parameters x_{t,K_i}^i to the server where
 126 it updates its global model, with $x_{t+1} = \sum_i \nu_i x_{t,K_i}^i / \nu$ where $\nu = \sum_i \nu_i$. To facilitate some of the
 127 discussions that follow, we define the following quantities:

$$\bar{\Delta}_t \triangleq x_{t+1} - x_t = \sum_{i=1}^m \frac{\nu_i}{\nu} \Delta_t^i \quad \text{where} \quad \Delta_t^i \triangleq - \sum_{k=0}^{K_i-1} \eta_{L,t} g_i(x_{t,k}^i) \quad (2)$$

128 **Offline Hyperparameter Tuning** Offline tuning is best to be summarized as follows. We first
 129 define $U = \{u \in \mathbb{R} \mid u \geq 0\}$ with $\eta_L \in U$, and $V = \{v \in \mathcal{I} \mid v \geq 1\}$ with $K \in V$. We also define
 130 $C = U \times V$, and $c = (\eta_L, K)$, where $c \in C$. Offline tuning would have the following objective:

131 $\min_{c \in C} f_{\text{valid}}(x, c)$ s.t. $x = \operatorname{argmin}_{z \in \mathbb{R}^d} f_{\text{train}}(z, c)$. With abuse of notation, we use f_{valid} for the
 132 objective function calculated from a validation dataset which is usually held-out before the procedure,
 133 and f_{train} for the objective from training data which usually is just local client data. A few notable
 134 offline tuning methods are as follows. Global grid-search from [Holly et al. \[2021\]](#) is an example
 135 of offline tuning that iterates over the entire search grid defined as C , completing an optimization
 136 process for each grid point and evaluating the result with a held-out validation set. Global Bayesian
 137 Optimization from [Holly et al. \[2021\]](#) is another similar example of offline tuning that follows the
 138 same template and objective. Instead of brute-force grid-search, c is sampled from a distribution \mathcal{D}_C
 139 over C , i.e. $c \sim \mathcal{D}_C$, that updates after every iteration.

140 **Online Hyperparameter Optimization** We are interested in an online procedure that combines
 141 hyperparameter optimization and model parameter optimization, with the following objective:

$$\min_{\substack{x \in \mathbb{R}^d \\ c \in C}} f_{\text{train}}(x, c) \quad (3)$$

142 This formulation is the objective of our method, FATHOM, which we will discuss shortly in detail. It
 143 has the advantage of joint optimization in a one-shot procedure. Furthermore, it does not assume the
 144 availability of a validation dataset.

145 3.3 Our Method: FATHOM

146 In this section we will introduce our method, FATHOM (Federated AuTomatic Hyperparameter
 147 OptiMization). Recall from our joint objective, eq(3), that both the model parameters, x , and
 148 hyperparameters of the optimization algorithm, c , are optimized jointly to minimize our objective
 149 function. An alternative view is to treat c as part of the parameters being optimized in a classic
 150 formulation, i.e. $\min_y f(y)$ with $y = (x, c)$. As previously mentioned, our method is inspired by
 151 hypergradient descent from [Baydin et al. \[2017\]](#) and by exponentiated gradient from [Amid et al.
 152 \[2022\]](#), both proposed for centralized learning rate optimization. We will present how FATHOM
 153 exploits our knowledge of analytical gradients to update client learning rate, number of local steps, as
 154 well as batch size, for an online, one-shot optimization procedure.

155 **Assumption 2.** (Convexity w.r.t. η_L and K) We assume $\mathbb{E}_t(f(x_t))$ is convex w.r.t. η_L and K , even
 156 though we assume non-convexity w.r.t. x_t). Specifically, convexity w.r.t. K follows the definition in
 157 [Murota \[1998\]](#), to accommodate the integer space where K is defined.

158 **Remark 1.** Assumption 2 is necessary to guarantee the existence of subgradients derived in Theorems
 159 1 and 2, and it will be assumed for this work. In problems dealing with deep neural networks, it is
 160 reasonable to not assume convexity w.r.t. hyperparameters. However, from our empirical results, we
 161 claim that the proposed algorithm is still able to operate as desired under this condition.

162 3.3.1 Hypergradient for Client Learning Rate

163 In this section, we derive the hypergradient for client learning rate in a similar fashion as [Baydin et al.
 164 \[2017\]](#), with the difference being that they are mainly concerned with the centralized optimization
 165 problem, and that we are concerned with the distributed setting where clients take local steps. We
 166 derive the following hypergradient of the objective function as defined in eq(1), taken with respect to
 167 the learning rate $\eta_{L,t-1}$ such that it can be updated to obtain $\eta_{L,t}$:

$$H_t = \frac{\partial f(x_t)}{\partial \eta_{L,t-1}} = \frac{\partial f(x_t)}{\partial x_t} \cdot \frac{\partial(x_{t-1} + \bar{\Delta}_{t-1})}{\partial \eta_{L,t-1}} = \nabla f(x_t) \cdot \frac{\partial \bar{\Delta}_{t-1}}{\partial \eta_{L,t-1}} \quad (4)$$

168 where $\bar{\Delta}_t$ is the update step for the global parameters x_t as defined in eq(2), leading to $\frac{\partial \bar{\Delta}_t}{\partial \eta_{L,t}} = \frac{\bar{\Delta}_t}{\eta_{L,t}} =$
 169 $-\sum_{i=1}^m \frac{\nu_i}{\nu} \sum_{k=0}^{K-1} g_i(x_t^{i,k})$. We also make the approximation $x_{t+1} - x_t = \bar{\Delta}_t \approx -\eta_{L,t} \nabla f(x_t)$. We
 170 can then write the normalized update, \bar{H}_t , similar to [Amid et al. \[2022\]](#), as follows:

$$\bar{H}_t = \frac{\nabla f(x_t)}{\|\nabla f(x_t)\|} \cdot \left(\frac{\partial \bar{\Delta}_{t-1}}{\partial \eta_{L,t-1}} / \left\| \frac{\partial \bar{\Delta}_{t-1}}{\partial \eta_{L,t-1}} \right\| \right) \approx -\frac{\bar{\Delta}_t}{\|\bar{\Delta}_t\|} \cdot \frac{\bar{\Delta}_{t-1}}{\|\bar{\Delta}_{t-1}\|} \quad (5)$$

171 The resulting hypergradient is a scalar, as expected, and can be used efficiently as part of the update
 172 rule for η_L , which we will see in Section 3.3.4. The implementation is communication efficient, since
 173 in each round, each client needs one extra scalar to send back to the server, and likewise the server
 174 needs to broadcast one extra scalar back to the clients. It is also computationally efficient since it
 175 avoids calculating the full local gradient $\nabla f(x_t)$.

176 3.3.2 Hypergradient for Number of Local Steps

177 Since the number of local steps is an integer, i.e. $K = \{k \in \mathbb{I} \mid k \geq 1\}$, this means $f(x_t)$ does not
 178 exist for non-integer values of K . We formulate a subgradient as a surrogate of the hypergradient
 179 $\partial f(x_t)/\partial K$, as follows. We will call this a hyper-subgradient.

180 **Theorem 1.** *When a piecewise function L_t is defined for every value of $K_0 \in [K]$ on l , such that
 181 $0.0 \leq l < 1.0$, we claim, under Assumption 2 that the following is a subgradient of $f(x_t)$ at
 182 $K_t = K_0$:*

$$\frac{\partial L_t}{\partial l} = \nabla f(x_t) \cdot \left(-\eta_{L,t} \sum_{i=1}^m g_i(x_{t-1}^{i,K_t-1}) \frac{\nu_i}{\nu} \right) \quad (6)$$

183 where l represents the marginal fraction of local steps beyond K_0 . We leave the proof (with an
 184 illustration in Figure 2) in the Appendix section beginning in eq(20).

185 The result from Theorem 1 is not sufficiently communication-efficient for implementing an update
 186 rule for K . This is because it would require the quantity $g_i(x_{t-1}^{i,K_t-1})$ to be communicated from
 187 each client i to the server. To save communication, let us reuse what the server has in memory:
 188 $\bar{\Delta}_t = \left(-\eta_L \sum_{i=1}^m \frac{\nu_i}{\nu} \sum_{k=0}^{K_t-1} g_i(x_t^{i,k}) \right)$. If we let:

$$S_t = \nabla f(x_t) \cdot \left(-\eta_{L,t} \sum_{i=1}^m \frac{\nu_i}{\nu} \sum_{k=0}^{K_t-1} g_i(x_{t-1}^{i,k}) \right) l \quad (7)$$

$$N_t = \frac{\partial S_t}{\partial l} = \nabla f(x_t) \cdot \left(-\eta_{L,t} \sum_{i=1}^m \frac{\nu_i}{\nu} \sum_{k=0}^{K_t-1} g_i(x_{t-1}^{i,k}) \right) = \nabla f(x_t) \cdot \bar{\Delta}_{t-1} \quad (8)$$

$$\bar{N}_t = \frac{\nabla f(x_t)}{\|\nabla f(x_t)\|} \cdot \frac{\bar{\Delta}_{t-1}}{\|\bar{\Delta}_{t-1}\|} \approx -\frac{\bar{\Delta}_t}{\|\bar{\Delta}_t\|} \cdot \frac{\bar{\Delta}_{t-1}}{\|\bar{\Delta}_{t-1}\|} \quad (9)$$

189 where eq(9) is the normalized update as in Amid et al. [2022]. We claim that eq(8) is a positively-
 190 biased version of eq(6), which has its practical importance due to the fact that the last term in eq(6)
 191 from Theorem 1 results in zero-mean, noisy gradients, when the local functions are nearing their local
 192 solutions, when in fact, this is the area where more local work is not needed. Thus, a positive bias is
 193 desirable to drive the number of local steps down. This result is also useful from a communication
 194 efficiency perspective in its implementation, because the server has all the components to calculate
 195 this quantity, and would not require additional communication.

196 3.3.3 Regularization for Number of Local Steps

197 One of the goals for FATHOM is savings in local computation. To avoid excessive number of local
 198 steps, we further develop a regularization term for local computation against excessive K , which is a
 199 proxy for the hypergradient of the local client functions at the end of each round : $\partial f_i(x_t^{i,K})/\partial K$.

200 **Theorem 2.** *When a piecewise function J_t is defined for every value of $K_0 \in [K]$ on l , such that
 201 $0.0 \leq l < 1.0$, we claim, under Assumption 2 that the following is a subgradient of $\sum_{i=1}^m f_i(x_t^{i,K_t})$
 202 at $K_t = K_0$:*

$$\frac{\partial J_t}{\partial l} = -\eta_{L,t} \sum_{i=1}^m \frac{\nu_i}{\nu} \mathbb{E} [g_i(x_t^{i,K_0-1})] \cdot g_i(x_t^{i,K_t}) \approx -\eta_{L,t} \sum_{i=1}^m \frac{\nu_i}{\nu} \sum_{k=0}^{K_t-1} g_i(x_t^{i,k}) \cdot g_i(x_t^{i,K_t}) \quad (10)$$

203 where l represents the marginal fraction of local steps beyond K_0 . We leave the proof in the Appendix
 204 section beginning in eq(24).

205 In our algorithm, we use the normalized update based on the following biased proxy, since eq(10)
 206 tends to be noisy from $g_i(x_t^{i,K_t})$.

$$G_t = -\eta_{L,t} \sum_{i=1}^m \frac{\nu_i}{\nu} \min_{K \leq K_t} \left(\sum_{k=0}^{K-1} g_i(x_t^{i,k}) \cdot g_i(x_t^{i,K}) \right) \quad (11)$$

$$\bar{G}_t = -\eta_{L,t} \sum_{i=1}^m \frac{\nu_i}{\nu} \min_{K \leq K_t} \left(\frac{\sum_{k=0}^{K-1} g_i(x_t^{i,k})}{\|\sum_{k=0}^{K-1} g_i(x_t^{i,k})\|} \cdot \frac{g_i(x_t^{i,K})}{\|g_i(x_t^{i,K})\|} \right) \quad (12)$$

207 where \bar{G}_t is the normalized update. The proxy yields a bias towards smaller number of local steps,
 208 which is desirable for reducing local computation. We use this biased proxy against using a more
 209 typical regularization such as L2 for the number of local steps, based on initial empirical results for
 210 better performance..

211 3.3.4 Normalized Exponentiated Gradient Updates

212 For the update rules of the hyperparameters η_L (client learning rate) and K (number of client local
 213 steps), we use the normalized exponentiated gradient descent method (EGN) with no momentum,
 214 rather than a conventional linear update method such as the additive update of hypergradient descent
 215 proposed in [Baydin et al. \[2017\]](#). It is reasonable to use exponentiated gradient (EG) methods for
 216 updates of hyperparameters that are strictly positive in value. EG methods also enjoy significantly
 217 faster convergence properties when only a small subset of the dimensions are relevant, according to
 218 [Amid et al. \[2022\]](#).

219 EG methods have been proposed in previous works for a variety of applications ([Khodak et al. \[2021\]](#),
 220 [Amid et al. \[2022\]](#), [Li et al. \[2020\]](#)), and analyzed in depth ([Ghai et al. \[2019\]](#)), where its convergence
 221 has been studied and validated ([Li and Cevher \[2018\]](#)). Recently, [Amid et al. \[2022\]](#) showed that EGN
 222 is the same as the multiplicative update for hypergradient descent proposed in [Baydin et al. \[2017\]](#),
 223 when the approximation $\exp(\cdot) \approx 1 + \cdot$ is made. From our observations, we believe that momentum
 224 is not needed for the effectiveness of EGN in our application, as validated in our numerical results.
 225 We also opted-out of adding further complexity such as extra weights and activation functions to
 226 model the relationships between $\eta_{L,t}$ and K_t , because it would require more samples to optimize and
 227 because FATHOM is a one-shot procedure. Furthermore, due to the non-stationary nature of these
 228 values, we opt for a simpler scheme for faster performance.

229 Hence, for the update rule of client learning rate, η_L , we have:

$$\eta_{L,t+1} = \eta_{L,t} \exp(-\gamma_\eta \bar{H}_t) \quad (13)$$

230 where \bar{H}_t is as defined in eq(5). For number of local steps, we observe that it is related to batch
 231 size in round t , B_t , as follows. To accommodate heterogeneity of local dataset sizes among clients,
 232 we have number of local data samples from client i to be ν_i . The number of local steps for client i
 233 is $K_i = \lfloor \nu_i E_t / B_t \rfloor$, where E_t is number of epochs, with $E_t = 1$ meaning the entire local dataset
 234 for each client to be processed once per round. We derive update rules for E_t and B_t globally to
 235 optimize the number of local steps, without having to make any changes to our theoretical analysis to
 236 accommodate the heterogeneity of local dataset sizes:

$$E_{t+1} = E_t \exp(-\gamma_E (\bar{N}_t + \bar{G}_t)) \quad (14)$$

237 and

$$B_{t+1} = B_t \exp(-\gamma_B (-\bar{G}_t)) \quad (15)$$

238 where N_t and G_t are defined in eq(9) and eq(12), respectively. These update rules accomplish the goal
 239 of updating the number of local steps via E_t/B_t with $\frac{E_{t+1}}{B_{t+1}} = \frac{E_t}{B_t} \exp(-\gamma_E \bar{N}_t - (\gamma_E - \gamma_B) \bar{G}_t)$.

240 Typically, with $\gamma_B \geq \gamma_E$, $(\gamma_B - \gamma_E) \bar{G}_t$ becomes a tunable regularization term as discussed at the
 241 end of Section 3.3.3

242 3.3.5 Client Sampling

243 We present our method, FATHOM, as shown in Algorithm 1. One practical factor we have not
 244 considered in our discussions is partial client sampling. For our implementation to handle the
 245 stochastic nature of client sampling, the metric $\bar{\Delta}_{t-1}$ for calculating \bar{H}_t in eq(5) and \bar{N}_t in eq(9) is
 246 modified by a smoothing function for noise filtering, i.e. $\bar{\Delta}_{t,sm} = \alpha \bar{\Delta}_{t-1,sm} + (1-\alpha) \bar{\Delta}_t$, which is a
 247 single-pole infinite impulse response filter ([Oppenheim and Schaffer \[2009\]](#) [Oppenheimer et al. \[2009\]](#))
 248 with no bias compensation. We use the notation "sm" for smoothed, and after many experiments, we
 249 decide to use $\alpha = 0.5$ for all of our numerical results.

Algorithm 1: FATHOM : $g_i(x)$ is defined in Assumptions [1](#) and m is the number of clients.

Input: Server initializes global model $x_{t=1}$, T as the end communication round, and:

$$\bar{\Delta}_{t=0,sm} = 0; \alpha = 0.5; \gamma_\eta = 0.01; \gamma_E = 0.01; \gamma_B = 0.1$$

Output: x_T , as well as $\eta_{L,t}$, E_t and B_t for all $t \in [T]$

for $t = 1, \dots, T$ **do**

 Sample client set S_t out of m clients.

 For each client $i \in S_t$, initialize: $x_t^{i,k=0} = x_t$ and $K_{t,i} = \lfloor \nu_i E_t / B_t \rfloor$.

 Set $\Delta_i = 0$, and $\phi_i = +\infty$.

for $k = 0, \dots, K_{t,i} - 1$ **do**

 For each client i , compute in parallel an unbiased stochastic gradient $g_i(x_t^{i,k})$.

 For each client i , calculate $\phi_i = \min(\phi_i, g_i(x_t^{i,k}) \cdot \Delta_i)$ where $\Delta_i = x_t^{i,k} - x_t$

 For each client i , update in parallel its local solution: $x_t^{i,k+1} = x_t^{i,k} - \eta_{L,t} g_i(x_t^{i,k})$

end

 Server calculates $\nu = \sum_{i \in S_t} \nu_i$, where ν_i is the size of client i dataset.

 Server calculates $\bar{\Delta}_t = \sum_{i \in S_t} \Delta_i (\nu_i / \nu)$; see eq[\(2\)](#)

 Server updates global model $x_{t+1} = x_t - \bar{\Delta}_t$

 Server calculates $\bar{H}_t = \bar{N}_t = -\frac{\bar{\Delta}_t}{\|\bar{\Delta}_t\|} \cdot \frac{\bar{\Delta}_{t-1,sm}}{\|\bar{\Delta}_{t-1,sm}\|}$, modified from eq[\(5\)](#) and eq[\(9\)](#)

 Server calculates \bar{G}_t ; see eq[\(12\)](#)

 Server updates client learning rate $\eta_{L,t+1}$, epochs, E_{t+1} , and batch size B_{t+1} for the next round; see eq[\(13\)](#), eq[\(14\)](#), and eq[\(15\)](#).

 Server updates $\Delta_{t,sm} = (1 - \alpha)\bar{\Delta}_t + \alpha\bar{\Delta}_{t-1,sm}$ for the next round

end

251 4 Theoretical Convergence

252 A standard approach to theoretical analysis of an online optimization method such as ours, is through
 253 analyzing the regret bound (Zinkevich [\[2003\]](#), Khodak et al. [\[2019\]](#), Kingma and Ba [\[2014\]](#), and
 254 Mokhtari et al. [\[2016\]](#)). Nonetheless, this approach does not tell us the impact on communication
 255 efficiency by the online updates introduced from FATHOM. Therefore, we take an alternative
 256 approach by extending the guarantees of FedAvg performance (Wang et al. [\[2021\]](#), Reddi et al.
 257 [\[2020\]](#), Gorbunov et al. [\[2020\]](#), Yang et al. [\[2021\]](#), Li et al. [\[2019\]](#), etc) to include both adaptive
 258 learning rate and adaptive number of local steps. We assume the special case in our analysis to have
 259 full client participation. We prove that adaptive learning rate and adaptive number of local steps does
 260 not impact asymptotic convergence, despite the given relaxed conditions.

261 4.1 Assumptions

262 **Assumption 3.** (*L-Lipschitz Continuous Gradient for Parameters x_t*) There exists a constant $L > 0$,
 263 such that $\|\nabla f_i(x) - \nabla f_i(y)\| \leq L\|x - y\|$, $\forall x, y \in \mathbb{R}^d$, and $i \in [m]$, where x and y are the
 264 parameters in eq[\(1\)](#)

265 **Assumption 4.** (*Bounded Local Variance*) There exist a constant $\sigma_L > 0$, such that the variance of
 266 each local gradient estimator is bounded by $\mathbb{E}\|\nabla f_i(x) - g_i(x)\|^2 \leq \sigma_L^2$, $\forall x$, and $i \in [m]$.

267 **Assumption 5.** (*Bounded Second Moment*) There exists a constant $G > 0$, such that $\mathbb{E}_t\|\nabla f_i(x_t)\| \leq$
 268 G , $i \in [m]$, $\forall x_t$.

269 4.2 Convergence Results

270 **Theorem 3.** Under Assumptions [1](#)[5](#) and with full client participation, when FATHOM as shown
 271 in Algorithm [1](#) is used to find a solution x_* to the unconstrained problem defined in eq[\(1\)](#), the
 272 sequence of outputs $\{x_t\}$ satisfies the following upper-bound, where, with slight abuse of notation,
 273 $\mathcal{E} = \min_{t \in [T]} \mathbb{E}_t \|\nabla f(x_t)\|_2^2$:

$$\mathcal{E}_{fathom} = \mathcal{O}\left(\sqrt{\frac{\sigma_L^2 + G^2}{mKT}} + \sqrt[3]{\frac{\sigma_L^2}{KT^2}} + \sqrt[3]{\frac{G^2}{T^2}}\right) \quad (16)$$

274 with the following conditions: $\bar{\eta}_L = \min \left(\sqrt{\frac{2\beta_0 m D}{\beta_1 \bar{K} L T (\sigma_L^2 + G^2)}}, \sqrt[3]{\frac{\beta_0 D}{2.5\beta_2 \bar{K}^2 L^2 \sigma_L^2 T}}, \sqrt[3]{\frac{\beta_0 D}{2.5\beta_3 \bar{K}^3 L^2 G^2 T}} \right)$

275 and $\eta_{L,t} \leq 1/L$ for all t , where

$$\bar{\eta}_L \triangleq \frac{1}{T} \sum_{t=1}^T \eta_{L,t} \quad \text{and} \quad \bar{K} \triangleq \frac{1}{T} \sum_{t=1}^T K_t \quad (17)$$

276 and where

$$\beta_0 = \frac{\sum_t \eta_{L,t} K_t}{T \left[\frac{1}{T} \sum_t \eta_{L,t} \right] \left[\frac{1}{T} \sum_t K_t \right]}, \quad \beta_1 = \frac{\sum_t \eta_{L,t} K_t \left[\frac{1}{T} \sum_t \eta_{L,t} \right]}{\sum_t \eta_{L,t}^2 K_t} \quad (18)$$

$$\beta_2 = \frac{\sum_t \eta_{L,t} K_t \left[\frac{1}{T} \sum_t \eta_{L,t} \right]^2 \left[\frac{1}{T} \sum_t K_t \right]}{\sum_t \eta_{L,t}^3 K_t^2}, \quad \beta_3 = \frac{\sum_t \eta_{L,t} K_t \left[\frac{1}{T} \sum_t \eta_{L,t} \right]^2 \left[\frac{1}{T} \sum_t K_t \right]^2}{\sum_t \eta_{L,t}^3 K_t^3} \quad (19)$$

277 We leave the proof in the Appendix beginning in eq(29).

278 The values of $\beta_0, \beta_1, \beta_2, \beta_3$, and β_4 are dependent on the relative changes over the adaptive process
 279 of these components, according to Chebyshev's Sum Inequalities (Hardy et al. [1988]). A special
 280 case is when these quantities equal to 1 when both $\eta_{L,t}$ and K_t are constant, which recovers the
 281 standard upperbound for FedAvg from eq(16).

282 **Remark 2.** The definitions in eq(17) combined with the conditions for $\bar{\eta}_L$ above is called the relaxed
 283 conditions in this paper for the hyperparameters $\eta_{L,t}$ and K_t . The values of $\eta_{L,t}$ and K_t are adaptive
 284 during the optimization process between rounds $t = 1$ and $t = T$, as long as the above conditions are
 285 satisfied for the guarantee in eq(31) to hold. This relaxation presents opportunities for a scheme such
 286 as FATHOM to exploit for performance gain. For example, suppose T approaches ∞ for a prolonged
 287 training session. Then $\bar{\eta}_L$ would necessarily be sufficiently small for \mathcal{E}_{fathom} to be bounded by
 288 eq(16). However, for early rounds i.e. small t values, $\eta_{L,t} \leq T\bar{\eta}_L$ can be reasonably large and still
 289 can satisfy eq(17), for the benefit of accelerated learning and convergence progress early on. Similar
 290 strategy can be used for number of local steps to minimize local computations towards later rounds.
 291 In any case, these strategies are mere guidelines meant to remain within the worst case guarantee.
 292 However, Theorem 3 offers the flexibility otherwise not available. We will now show the empirical
 293 performance gained by taking advantage of this flexibility.

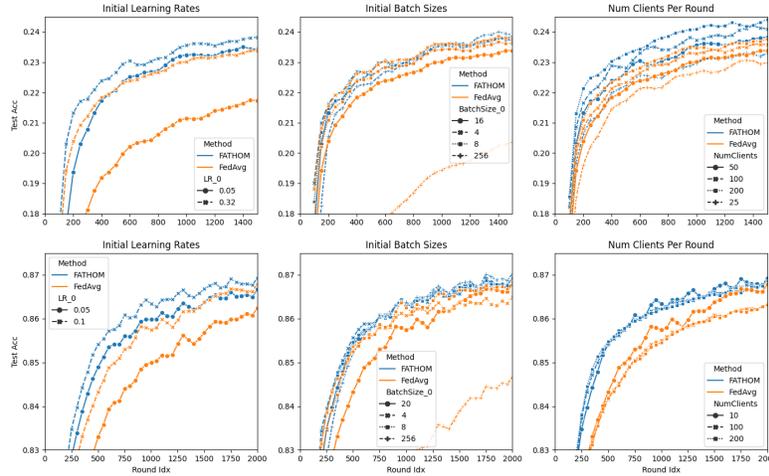


Figure 1: Test Accuracy Performance with various values of initial client learning rate (LR_0), initial batch size (BatchSize_0), and number of clients per round (NumClients). Top row: FSO sims. Bottom row: FEMNIST sims. Baseline values for FEMNIST: LR_0=0.1, BatchSize_0=20, NumClients=10. Baseline values for FSO: LR_0=0.32, BatchSize_0=16, NumClients=50.

294 5 Empirical Evaluation and Numerical Results

295 We present an empirical evaluation of FATHOM proposed in Section 3 and outlined in Algorithm
 296 1. We conduct extensive simulations of federated learning in character recognition on the federated
 297 EMNIST-62 dataset (FEMNIST) (Cohen et al. [2017]) with a CNN, and in natural language next-word
 298 prediction on the federated Stack Overflow dataset (FSO) (TensorFlow-Federated-Authors [2019])
 299 with a RNN. We defer most of the details of the experiment setup in Appendix Section C.1. Our
 300 choice of datasets, tasks and models, are exactly the same as the "EMNIST CR" task and the "SO
 301 NWP" task from Reddi et al. [2020]. See Figure 1 and Table 1 and their captions for details of the
 302 experiment results. Our evaluation lacks comparison with a few one-shot FL HPO methods discussed
 303 earlier in the paper because of a lack of standardized benchmark (until FedHPO- B Wang et al. [2022])
 304 was published concurrently as this work) to be fair and comprehensive.

305 The underlying principle behind these experiments is evaluating the robustness of FATHOM versus
 306 FedAvg under various initial settings, to mirror realistic usage scenarios where the optimal hyperpa-
 307 rameter values are unknown. For FATHOM, we start with the same initial hyperparameter values
 308 as FedAvg. The test accuracy progress with respect to communication rounds is shown in Figure 1
 309 from these experiments. We also pick test accuracy targets for the two tasks. For FEMNIST CR we
 310 use 86% and for FSO NWP we use 23%. Table 1 shows a table of resource utilization metrics with
 311 respect to reaching these targets in our experiments, highlighting the communication efficiency as
 312 well as reduction in local computation from FATHOM in comparison to FedAvg. To our knowledge,
 313 we are the first to show gain from an online HPO procedure over a well-tuned equivalent procedure
 314 with fixed hyperparameter values.

315 The federated learning simulation framework on which we build our algorithms for our experiments
 316 is FedJAX (Ro et al. [2021]) which is under the Apache License. The server that runs the experiments
 317 is equipped with Nvidia Tesla V100 SXM2 GPUs.

Table 1: Resource utilization in communication and local computation to reach specified test accuracy target for each task. All evaluations are run for ten trials. Bold numbers highlight better performance. NA means target was not reached within 1500 rounds for FSO NWP and 2000 rounds for FEMNIST CR, in any of our trials. LR_0 is initial client learning rate, BS_0 is initial batch size, and NCPR is number of clients per round. All experiments use baseline initial values except where indicated. For clarification, M is used in place for "million", and K for "thousand".

Baseline_fso : (LR_0 = 0.32, BS_0 = 16, NCPR = 50)

Baseline_femnist : (LR_0 = 0.10, BS_0 = 20, NCPR = 10)

Tasks	Experiments	Number of Rounds To Reach Target Test Accuracy		Local Gradients Calculated To Reach Target Test Accuracy	
		FATHOM	FedAvg	FATHOM	FedAvg
FSO NWP Target@23%	Baseline_fso	562 ± 12	971 ± 11	85M ± 1.2M	124M ± 1.3M
	LR_0 = 0.05	871 ± 7	NA	138M ± 3.2M	NA
	BS_0 = 4	758 ± 43	580 ± 18	93M ± 2.8M	74M ± 2.5M
	BS_0 = 256	801 ± 28	NA	174M ± 18M	NA
	NCPR = 25	970 ± 49	1283 ± 33	63M ± 2.7M	82M ± 3.8M
	NCPR = 200	396 ± 17	684 ± 26	280M ± 45M	350M ± 13M
FEMNIST CR Target@86%	Baseline_femnist	739 ± 24	1098 ± 15	1.5M ± 36K	2.2M ± 64K
	LR_0 = 0.05	905 ± 21	1574 ± 19	1.7M ± 28K	3.1M ± 28K
	BS_0 = 4	708 ± 17	885 ± 41	1.2M ± 28K	1.7M ± 88K
	BS_0 = 256	736 ± 20	NA	2.0M ± 44K	NA
	NCPR = 100	777 ± 16	1436 ± 18	22M ± 0.27M	28M ± 0.39K
	NCPR = 200	790 ± 16	1481 ± 33	57M ± 1.0M	59M ± 1.3M

318 6 Conclusion and Future Work

319 In this work, we propose FATHOM for adaptive hyperparameters in federated optimization, specifi-
 320 cally for FedAvg. We analyze theoretically and evaluate empirically its potential benefits in conver-
 321 gence behavior as measured in test accuracy, and in reduction of local computations, by automatically
 322 adapting the three main hyperparameters of FedAvg: client learning rate, and number of local steps
 323 via epochs and batch size. An example of future efforts to extend this work is using a standardized
 324 benchmark such as Wang et al. [2022] for performance comparison against other FL HPO methods.

325 References

- 326 E. Amid, R. Anil, C. Fifty, and M. K. Warmuth. Step-size adaptation using exponentiated gradient updates,
327 2022.
- 328 A. G. Baydin, R. Cornish, D. M. Rubio, M. Schmidt, and F. Wood. Online learning rate adaptation with
329 hypergradient descent, 2017.
- 330 H. Cai, L. Zhu, and S. Han. ProxylesNAS: Direct neural architecture search on target task and hardware.
331 In *International Conference on Learning Representations*, 2019. URL <https://arxiv.org/pdf/1812>
332 [00332.pdf](https://arxiv.org/pdf/181200332.pdf)
- 333 Z. Charles and J. Konečný. On the outsized importance of learning rates in local update methods, 2020.
- 334 G. Cohen, S. Afshar, J. Tapson, and A. Van Schaik. Emnist: Extending mnist to handwritten letters. In *2017*
335 *international joint conference on neural networks (IJCNN)*, pages 2921–2926. IEEE, 2017.
- 336 Z. Dai, K. H. Low, and P. Jaillet. Federated bayesian optimization via thompson sampling, 2020.
- 337 Z. Dai, B. K. H. Low, and P. Jaillet. Differentially private federated bayesian optimization with distributed
338 exploration, 2021.
- 339 U. Ghai, E. Hazan, and Y. Singer. Exponentiated gradient meets gradient descent, 2019.
- 340 E. Gorbunov, F. Hanzely, and P. Richtárik. Local sgd: Unified theory and new efficient methods, 2020.
- 341 P. Guo, D. Yang, A. Hatamizadeh, A. Xu, Z. Xu, W. Li, C. Zhao, D. Xu, S. Harmon, E. Turkbey, et al. Auto-fedrl:
342 Federated hyperparameter optimization for multi-institutional medical image segmentation. *arXiv preprint*
343 *arXiv:2203.06338*, 2022.
- 344 G. Hardy, J. Littlewood, and G. Pólya. *Inequalities*. Cambridge Mathematical Library. Cambridge University
345 Press, 1988. ISBN 9781107647398. URL <https://books.google.com/books?id=EfvZAQAQBAJ>
- 346 S. Holly, T. Hiessl, S. R. Lakani, D. Schall, C. Heitzinger, and J. Kemnitz. Evaluation of hyperparameter-
347 optimization approaches in an industrial federated learning system, 2021.
- 348 M. Khodak, M.-F. Balcan, and A. Talwalkar. Adaptive gradient-based meta-learning methods, 2019.
- 349 M. Khodak, R. Tu, T. Li, L. Li, N. Balcan, V. Smith, and A. Talwalkar. Federated hyperparameter tuning:
350 Challenges, baselines, and connections to weight-sharing. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. W.
351 Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL [https://openreview](https://openreview.net/forum?id=p99rWde9fVJ)
352 [net/forum?id=p99rWde9fVJ](https://openreview.net/forum?id=p99rWde9fVJ)
- 353 D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2014.
- 354 D. P. Kingma and M. Welling. Auto-encoding variational bayes, 2013.
- 355 L. Li, M. Khodak, M.-F. Balcan, and A. Talwalkar. Geometry-aware gradient algorithms for neural architecture
356 search, 2020.
- 357 X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang. On the convergence of fedavg on non-iid data, 2019.
- 358 Y.-H. Li and V. Cevher. Convergence of the exponentiated gradient method with armijo line search. *Journal*
359 *of Optimization Theory and Applications*, 181(2):588–607, Dec 2018. ISSN 1573-2878. doi: 10.1007/
360 s10957-018-1428-9. URL <http://dx.doi.org/10.1007/s10957-018-1428-9>
- 361 H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of
362 deep networks from decentralized data, 2016.
- 363 A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro. Online optimization in dynamic environments:
364 Improved regret rates for strongly convex problems. *2016 IEEE 55th Conference on Decision and Control*
365 *(CDC)*, Dec 2016. doi: 10.1109/cdc.2016.7799379. URL <http://dx.doi.org/10.1109/cdc.2016>
366 [7799379](http://dx.doi.org/10.1109/cdc.20167799379)
- 367 H. Mostafa. Robust federated learning through representation matching and adaptive hyper-parameters, 2019.
- 368 K. Murota. Discrete convex analysis. *Mathematical Programming*, 83:313–371, 1998.
- 369 A. V. Oppenheim and R. W. Schaffer. *Discrete-Time Signal Processing*. Prentice Hall Press, USA, 3rd edition,
370 2009. ISBN 0131988425.

- 371 H. Pham, M. Guan, B. Zoph, Q. Le, and J. Dean. Efficient neural architecture search via parameters sharing.
 372 In J. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*,
 373 volume 80 of *Proceedings of Machine Learning Research*, pages 4095–4104. PMLR, 10–15 Jul 2018. URL
 374 <https://proceedings.mlr.press/v80/pham18a.html>.
- 375 S. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, S. Kumar, and H. B. McMahan. Adaptive
 376 federated optimization, 2020.
- 377 J. H. Ro, A. T. Suresh, and K. Wu. Fedjax: Federated learning simulation with jax. *arXiv preprint*
 378 *arXiv:2108.02117*, 2021.
- 379 TensorFlow-Federated-Authors. Tensorflow federated stack overflow dataset, 2019. URL [https://www.
 380 tensorflow.org/federated/api_docs/python/tff/simulation/datasets/stackoverflow](https://www.tensorflow.org/federated/api_docs/python/tff/simulation/datasets/stackoverflow).
- 381 J. Wang and G. Joshi. Adaptive communication strategies to achieve the best error-runtime trade-off in local-
 382 update sgd, 2018.
- 383 J. Wang, Z. Charles, Z. Xu, G. Joshi, H. B. McMahan, B. A. y Arcas, M. Al-Shedivat, G. Andrew, S. Avestimehr,
 384 K. Daly, D. Data, S. Diggavi, H. Eichner, A. Gadhikar, Z. Garrett, A. M. Girgis, F. Hanzely, A. Hard, C. He,
 385 S. Horvath, Z. Huo, A. Ingerman, M. Jaggi, T. Javidi, P. Kairouz, S. Kale, S. P. Karimireddy, J. Konecny,
 386 S. Koyejo, T. Li, L. Liu, M. Mohri, H. Qi, S. J. Reddi, P. Richtarik, K. Singhal, V. Smith, M. Soltanolkotabi,
 387 W. Song, A. T. Suresh, S. U. Stich, A. Talwalkar, H. Wang, B. Woodworth, S. Wu, F. X. Yu, H. Yuan,
 388 M. Zaheer, M. Zhang, T. Zhang, C. Zheng, C. Zhu, and W. Zhu. A field guide to federated optimization, 2021.
- 389 Z. Wang, W. Kuang, C. Zhang, B. Ding, and Y. Li. Fedhpo-b: A benchmark suite for federated hyperparameter
 390 optimization, 2022.
- 391 R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning.
 392 *Mach. Learn.*, 8(3–4):229–256, may 1992. ISSN 0885-6125. doi: 10.1007/BF00992696. URL [https:
 393 //doi.org/10.1007/BF00992696](https://doi.org/10.1007/BF00992696).
- 394 H. Yang, M. Fang, and J. Liu. Achieving linear speedup with partial worker participation in non-iid federated
 395 learning, 2021.
- 396 Y. Zhou, P. Ram, T. Salonidis, N. Baracaldo, H. Samulowitz, and H. Ludwig. Single-shot hyper-parameter
 397 optimization for federated learning: A general algorithm and analysis, 2022.
- 398 M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of*
 399 *the Twentieth International Conference on International Conference on Machine Learning*, ICML’03, page
 400 928–935. AAAI Press, 2003. ISBN 1577351894.

401 Checklist

402 The checklist follows the references. Please read the checklist guidelines carefully for information on
 403 how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or
 404 **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing
 405 the appropriate section of your paper or providing a brief inline description. For example:

- 406 • Did you include the license to the code and datasets? **[Yes]** The code is MIT licensed.

407 Please do not modify the questions and only use the provided macros for your answers. Note that the
 408 Checklist section does not count towards the page limit. In your paper, please delete this instructions
 409 block and only keep the Checklist section heading above along with the questions/answers below.

410 1. For all authors...

- 411 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
 412 contributions and scope? **[Yes]**
- 413 (b) Did you describe the limitations of your work? **[Yes]** Please refer to Sections **2** and **6**
- 414 (c) Did you discuss any potential negative societal impacts of your work? **[No]** Not
 415 specifically, but it is alluded to how FL applications have social implications in the
 416 introductory section.
- 417 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
 418 them? **[Yes]**

- 419 2. If you are including theoretical results...
- 420 (a) Did you state the full set of assumptions of all theoretical results? [Yes] Please refer to
- 421 Assumptions [1](#), [2](#), [3](#), [4](#), and [5](#)
- 422 (b) Did you include complete proofs of all theoretical results? [Yes] Yes, in the supple-
- 423 mental material.
- 424 3. If you ran experiments...
- 425 (a) Did you include the code, data, and instructions needed to reproduce the main exper-
- 426 imental results (either in the supplemental material or as a URL)? [Yes] Yes, in the
- 427 supplemental material.
- 428 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
- 429 were chosen)? [Yes] Yes, in the supplemental material.
- 430 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
- 431 ments multiple times)? [Yes] Yes, see Table [1](#)
- 432 (d) Did you include the total amount of compute and the type of resources used (e.g., type
- 433 of GPUs, internal cluster, or cloud provider)? [Yes] Yes, it is mentioned in Section [5](#)
- 434 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 435 (a) If your work uses existing assets, did you cite the creators? [Yes] Yes, it is mentioned
- 436 in Section [5](#)
- 437 (b) Did you mention the license of the assets? [Yes] Yes, it is mentioned in Section [5](#)
- 438 (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
- 439 Yes, in the supplemental material.
- 440 (d) Did you discuss whether and how consent was obtained from people whose data you're
- 441 using/curating? [N/A]
- 442 (e) Did you discuss whether the data you are using/curating contains personally identifiable
- 443 information or offensive content? [No]
- 444 5. If you used crowdsourcing or conducted research with human subjects...
- 445 (a) Did you include the full text of instructions given to participants and screenshots, if
- 446 applicable? [N/A]
- 447 (b) Did you describe any potential participant risks, with links to Institutional Review
- 448 Board (IRB) approvals, if applicable? [N/A]
- 449 (c) Did you include the estimated hourly wage paid to participants and the total amount
- 450 spent on participant compensation? [N/A]