

FedHB: Hierarchical Bayesian Federated Learning

Minyoung Kim¹

¹*Samsung AI Center Cambridge
Cambridge, CB1 2JH, UK*

MIKIM21@GMAIL.COM

Timothy Hospedales^{1,2}

²*School of Informatics
The University of Edinburgh
Edinburgh, EH8 9AB, UK*

T.HOSPEDALES@ED.AC.UK

Editor: Matthew Hoffman

Abstract

We propose a novel hierarchical Bayesian approach to Federated Learning (FL), where our model reasonably describes the generative process of clients' local data via hierarchical Bayesian modeling: constituting random variables of local models for clients that are governed by a higher-level global variate. Interestingly, the variational inference in our Bayesian model leads to an optimisation problem whose block-coordinate descent solution becomes a distributed algorithm that is separable over clients and allows them not to reveal their own private data at all, thus fully compatible with FL. We also highlight that our block-coordinate algorithm has particular forms that subsume the well-known FL algorithms including Fed-Avg and Fed-Prox as special cases. Beyond introducing novel modeling and derivations, we also offer convergence analysis showing that our block-coordinate FL algorithm converges to an (local) optimum of the objective at the rate of $O(1/\sqrt{t})$, the same rate as regular (centralised) SGD, as well as the generalisation error analysis where we prove that the test error of our model on unseen data is guaranteed to vanish as we increase the training data size, thus asymptotically optimal.

Keywords: Federated Learning, Bayesian Methods, Probabilistic Models, Variational Inference, Block Coordinated Optimisation

1. Introduction

Federated Learning (FL) aims to enable a set of clients to collaboratively train a model in a privacy preserving manner, without sharing data with each other or a central server. Compared to conventional centralised optimisation problems, FL comes with a host of statistical and systems challenges – such as communication bottlenecks and sporadic participation. The key statistical challenge is non-i.i.d. data distributions across clients, each of which has a different data collection bias and potentially a different data labeling function (e.g., user preference learning). The classic and most popularly deployed FL algorithms are FedAvg (McMahan et al., 2017) and FedProx (Li et al., 2018), however, even when a global model can be learned, it often underperforms on each client's local data distribution in scenarios of high heterogeneity (Li et al., 2019; Karimireddy et al., 2019; Wang et al., 2020). Studies attempted to alleviate this by personalising learning at each client, allowing each local model to deviate from the shared global model (Sun et al., 2021). However, this

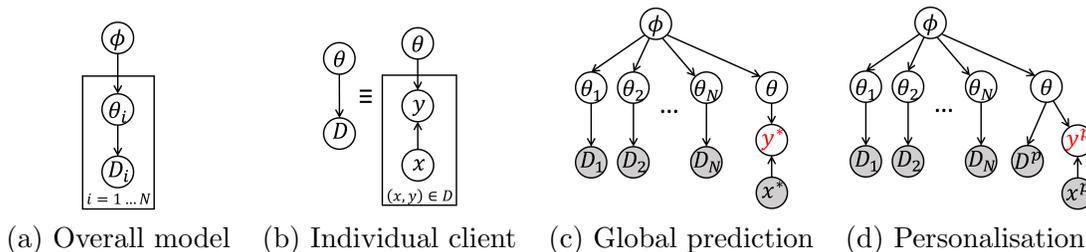


Figure 1: Graphical models. (a) Plate view of iid clients. (b) Individual client data with input images x given and only $p(y|x)$ modeled. (c) & (d): Global prediction and personalisation as probabilistic inference problems (shaded nodes = *evidences*, red colored nodes = *targets* to infer, x^* = test input in global prediction, D^p = training data for personalisation and x^p = test input).

remains challenging given that each client may have a limited amount of local data for personalised learning.

These challenges have motivated several attempts to model the FL problem from a Bayesian perspective. Introducing distributions on model parameters θ has enabled various schemes for estimating a global model posterior $p(\theta|D_{1:N})$ from clients’ local posteriors $p(\theta|D_i)$, or to regularise the learning of local models given a prior defined by the global model (Zhang et al., 2022; Al-Shedivat et al., 2021; Chen and Chao, 2021). However, these methods are not complete and principled solutions – having not yet provided full Bayesian descriptions of the FL problem, and having had resort to ad-hoc treatments to achieve tractable learning. The key difference is that they fundamentally treat network weights θ as a random variable shared across all clients. We introduce a *hierarchical* Bayesian model that assigns each client it’s own random variable for model weights θ_i , and these are linked via a higher level random variable ϕ as $p(\theta_{1:N}, \phi) = p(\phi) \prod_{i=1}^N p(\theta_i|\phi)$. This has several crucial benefits: Firstly, given this hierarchy, variational inference in our framework decomposes into separable optimisation problems over θ_i s and ϕ , enabling a practical Bayesian learning algorithm to be derived that is fully compatible with FL constraints, without resorting to ad-hoc treatments or strong assumptions. Secondly, this framework can be instantiated with different assumptions on $p(\theta_i|\phi)$ to deal elegantly and robustly with different kinds of statistical heterogeneity, as well as for principled and effective model personalisation.

Our resulting algorithm, termed Federated Hierarchical Bayes (FedHB) is empirically effective, as we demonstrate in a wide range of experiments on established benchmarks. More importantly, it benefits from rigorous theoretical support. In particular, we provide convergence guarantees showing that FedHB has the same $O(1/\sqrt{T})$ convergence rate as centralised SGD algorithms, which are not provided by related prior art (Zhang et al., 2022; Chen and Chao, 2021). We also provide a generalisation bound showing that FedHB is asymptotically optimal, which has not been shown by prior work such as (Al-Shedivat et al., 2021). Furthermore we show that FedHB subsumes classic methods FedAvg (McMahan et al., 2017) and FedProx (Li et al., 2018) as special cases, and ultimately provides additional justification and explanation for these seminal methods.

2. Related Work and Our Contributions

2.1 Related Work

General FL approaches. Perhaps the seminal pioneering work on FL is attributed to FedAvg (McMahan et al., 2017), which proposed fairly intuitive local training and global aggregation strategies with minimal training and communication complexity. A potential issue of divergence of global and local models due to the separated steps of local training and aggregation was addressed by model regularisation in the follow-up works (Li et al., 2018; Acar et al., 2021), which is shown to help the global model converge more reliably. Recent approaches aimed to incorporate benefits from existing machine learning approaches including domain adaptation/generalisation, clustering, multi-task learning, transfer learning, and meta-learning. To deal with heterogeneous client data distributions, those works in (Peterson et al., 2019; Zhang et al., 2021; Sun et al., 2021) attempted to tackle the FL problem in the perspective of (multi-)Domain Adaptation/Generalisation. Another interesting line of works aims to cluster clients with similar data distributions together (Briggs et al., 2020; Mansour et al., 2020). Along the line, the shared representations among the related or similar clients can be modeled motivated from general multi-task learning (Smith et al., 2017; Dinh et al., 2021). Motivated from transfer learning, reasonable attempts are made to exploit the idea of learning/transferring knowledge from related clients (Chen et al., 2020; Yang et al., 2020; Dinh et al., 2020; Li et al., 2021). There have been attempts to the personalised FL methods based on meta learning since the finetuning from the global trained model can be seen as adaptation to new data (Chen et al., 2018; Fallah et al., 2020).

Comparison to existing Bayesian FL approaches. Some recent studies tried to address the FL problem using Bayesian methods. As we mentioned earlier, the key difference is that these methods do not introduce Bayesian *hierarchy*, and ultimately treat network weights θ as a random variable *shared* across all clients, while our approach assigns individual θ_i to each client i governed by a common prior $p(\theta_i|\phi)$. The non-hierarchical approaches must all resort to ad hoc heuristics or strong assumptions in parts of their algorithm. More specifically, **FedPA** (Posterior Averaging) (Al-Shedivat et al., 2021) aims to establish the decomposition, $p(\theta|D_{1:N}) \propto \prod_{i=1}^N p(\theta|D_i)$ also known as *product of experts*, to allow client-wise inference/optimisation of $p(\theta|D_i)$. Unfortunately this decomposition does not hold in general unless we make a strong assumption of uninformative prior $p(\theta) \propto 1$ as they did. **FedBE** (Bayesian Ensemble) (Chen and Chao, 2021) aims to build the global posterior distribution $p(\theta|D_{1:N})$ from the individual posteriors $p(\theta|D_i)$ in either of two ad-hoc ways: SWAG (Maddox et al., 2019)-like model averaging over clients, or a convex combination of the modes of the local posteriors. **pFedBayes** (Zhang et al., 2022) can be seen as an implicit regularisation-based method to approximate $p(\theta|D_{1:N})$ from individual posteriors $p(\theta|D_i)$. To combine the individual posteriors, they introduce the so-called global distribution $w(\theta)$, which essentially serves as a *regulariser* that aims to enforce local posteriors $p(\theta|D_i)$ not to deviate from it, i.e., $p(\theta|D_i) \approx w(\theta)$ for all i . The introduction of $w(\theta)$ and its update strategy appears to be a hybrid treatment rather than solely Bayesian perspective. **FedEM** (Marfoq et al., 2021) forms a seemingly reasonable hypothesis that local client data distributions can be identified as mixtures of a fixed number of base distributions (with different mixing proportions). However, although they have probabilistic modeling, mixture estimation, and base distribution learning under this hypothesis, this method is not a Bayesian approach.

FedGP (Achituve et al., 2021b) aims to extend the GP-Tree algorithm (Achituve et al., 2021a) to the FL setting via the shared deep kernel learning. To this end, the clients perform GP-Tree kernel learning locally on its own data while the server aggregation simply follows the FedAvg algorithm to learn a global kernel. In this sense, the overall approach is quite different from our hierarchical Bayesian treatment. **FedPop** (Kotelevskii et al., 2022): It has a similar hierarchical Bayesian model structure as ours. But they split the backbone network parameters into those of the feature extractor (denoted by ϕ in the paper) and the linear classification head (β). In their model, the feature extractor weights ϕ are shared across the clients (called *fixed effects*), and the client-wise classification head parameters z^i are sampled from β , i.e., $z^i \sim p(z|\beta)$. Thus the client data D_i is generated by ϕ and z^i . The main differences from our approach are in four folds: 1) The higher-level variables β and local variables z^i sampled from β are both restricted to the linear classification head part of the network, which makes imposing uncertainty in model parameters quite limited; 2) Moreover, they do not actually treat β (and ϕ of feature extractor) as random variables, but deterministic variables which are optimized in empirical Bayes learning. This hinders the model from benefiting from hierarchical Bayesian modeling (e.g., they do not have prior distribution $p(\beta)$ at all); 3) Their optimization is alternating between the feature extractor ϕ and the head prior parameters β , utterly different from our block coordinate optimization alternating between higher level random variables and individual local variables; 4) They did not use variational inference for inference $p(z^i|D_i, \phi, \beta)$, but MCMC sampling (Langevin dynamics), which is the very reason why they had to reduce the size of the latents z^i only limited to classification heads, instead of full network parameters as we did.

Other Bayesian FL algorithms. Some approaches (Rendell et al., 2021; V. et al., 2021) proposed hierarchical Bayesian models for distributed inference that are similar to our model in (graphical model) structures. However, they used MCMC-based inference with Metropolis-Gibbs or Langevin-MC, which is the main difference from our approach that uses variational inference with block-coordinate descent optimisation. Practically, due to the computational overhead of MCMC, these methods were only tested for small data/networks (e.g., linear models or a single hidden-layer MLP). On the other hand, our approach is computationally more efficient, scalable to much larger backbones like MobileNet – the network size of their MLP is about 100K parameters while our MobileNet has 3.3M parameters, thus at $\times 30$ the scale. The work in (Louizos et al., 2021) proposed a similar hierarchical Bayesian approach to FL, but used the hard-EM algorithm to learn the higher-level parameters. Some recent Bayesian methods adopt the expectation-propagation (EP) approximations (Ashman et al., 2022; Guo et al., 2023): Unlike our model where each client has its own local model parameters, the former model only has a single random variable (model parameters) shared across clients. Their Partitioned variational inference in this work refers to the EP update steps that are done locally with the client data. In the latter work (Guo et al., 2023), there are no individual/local models for client data while they followed FedPA-like modeling $p(\theta|D_{1:N})$, thus an extension of FedPA. The use of EP is similar to the former where the EP update is done locally with the client data. Thus neither of these two works is a hierarchical Bayesian model, meaning that they lack a systematic way to distinctly model global and local parameters for global prediction and personalised prediction respectively. A hierarchical Bayesian approach with Normal-Inverse-Gamma (NIG) was previously introduced in (Fang et al., 2020). The general idea looks similar

to ours, but there are several key differences: (i) They used an EM algorithm that is not fully Bayesian; (ii) Their EM algorithm heavily relies on the Gaussian-NIG conjugacy, thus limiting the model class to linear models and unable to deal with deep nets; (iii) They only deal with global prediction (no personalisation); (iv) Our HB framework is very general, not limited to NIW; (v) Unlike their work, we also provide theoretical analysis on convergence and generalization error. In (Ozkara et al., 2023), a Gaussian mixture prior was adopted for personalisation, however, as a non-Bayesian regulariser in the deterministic optimisation.

2.2 Outline of Our Contributions

Our main contributions are summarised as follows:

1. To the best of our knowledge, our proposed approach is the first to show that variational hierarchical Bayesian inference formulation and its block-coordinate optimisation lead to a distributed algorithm that is fully compatible with the FL constraints.
2. We show that many well-known existing FL algorithms (e.g., Fed-Avg, Fed-Prox) can be subsumed as special cases in our general framework. Our Bayesian formulation generalises these algorithms by incorporating uncertainty (e.g., random dropout components), which is known to help regularise a model and lead to better generalisation.
3. Whereas many existing FL algorithms focus on only one of the two major FL tasks, namely global prediction and personalisation, we tackle both problems in a unifying principled Bayesian inference perspective.
4. We offer both convergence and generalisation error analysis for our general FL framework, showing that our proposed training algorithm converges to an (local) optimum at the same rate as the centralised SGD, with asymptotically optimal test performance.
5. Compared to existing hierarchical or flat Bayesian models, our approach is a general framework that can incorporate any prior/posterior distribution family without making specific distributional assumptions. Whereas many Bayesian methods are restricted in that only a small part (e.g., readout heads) of the deep nets are treated as random variables due to computational overhead, our approach is an efficient block-coordinate optimiser, supporting full Bayesian treatment of all parameters in deep nets.

3. Bayesian FL: General Framework

We introduce two types of latent random variables, ϕ and $\{\theta_i\}_{i=1}^N$. Each θ_i is deployed as the network weights for client i 's backbone. The variable ϕ can be viewed as a globally shared variable that is responsible for linking the individual client parameters θ_i . We assume conditionally independent and identical priors, $p(\theta_{1:N}|\phi) = \prod_{i=1}^N p(\theta_i|\phi)$. The prior for the latent variables $(\phi, \{\theta_i\}_{i=1}^N)$ is formed in a hierarchical fashion as (1). The local data for client i , denoted by D_i , is generated¹ by θ_i ,

$$\text{(Prior)} \quad p(\phi, \theta_{1:N}) = p(\phi) \prod_{i=1}^N p(\theta_i|\phi) \quad \text{(Likelihood)} \quad p(D_i|\theta_i) = \prod_{(x,y) \in D_i} p(y|x, \theta_i), \quad (1)$$

1. Note that we do not deal with generative modeling of input images x . Inputs x are always given, and only conditionals $p(y|x)$ are modeled. See Fig. 1(b) for the in-depth graphical model diagram.

where $p(y|x, \theta_i)$ is a conventional neural network model (e.g., softmax link for classification tasks). See the graphical model in Fig. 1(a) where the iid clients are governed by a single random variable ϕ .

Given the data D_1, \dots, D_N (also denoted by $D_{1:N}$), we infer the posterior, $p(\phi, \theta_{1:N}|D_{1:N}) \propto p(\phi) \prod_{i=1}^N p(\theta_i|\phi)p(D_i|\theta_i)$, which is intractable in general, and we adopt the variational inference to approximate it:

$$q(\phi, \theta_{1:N}; L) := q(\phi; L_0) \prod_{i=1}^N q_i(\theta_i; L_i), \quad (2)$$

where the variational parameters L consists of L_0 (parameters for $q(\phi)$) and $\{L_i\}_{i=1}^N$'s (parameters for $q_i(\theta_i)$'s from individual clients). Note that although θ_i 's are independent across clients under (2), they are differently modeled (the subscript i in notation q_i), reflecting different posterior beliefs originating from heterogeneity of local data D_i 's.

3.1 From Variational Inference to Federated Learning Algorithm

Using the standard variational inference techniques (Blei et al., 2017; Kingma and Welling, 2014), we can derive the ELBO objective function (details in Appendix A.1). We denote the *negative* ELBO by \mathcal{L} (to be minimised over L):

$$\mathcal{L}(L) := \sum_{i=1}^N \left(\mathbb{E}_{q_i(\theta_i)}[-\log p(D_i|\theta_i)] + \mathbb{E}_{q(\phi)}[\text{KL}(q_i(\theta_i)||p(\theta_i|\phi))] \right) + \text{KL}(q(\phi)||p(\phi)), \quad (3)$$

where we drop the dependency on L in notation for simplicity. Instead of optimizing (3) over the parameters L jointly as usual practice, we consider block-wise optimisation, also known as *block-coordinate optimisation* (Wright, 2015), specifically alternating two steps: (i) updating/optimizing all L_i 's $i = 1, \dots, N$ while fixing L_0 , and (ii) updating L_0 with all L_i 's fixed. That is,

- Optimisation over L_1, \dots, L_N (L_0 fixed).

$$\min_{\{L_i\}_{i=1}^N} \sum_{i=1}^N \left(\mathbb{E}_{q_i(\theta_i; L_i)}[-\log p(D_i|\theta_i)] + \mathbb{E}_{q(\phi; L_0)}[\text{KL}(q_i(\theta_i; L_i)||p(\theta_i|\phi))] \right). \quad (4)$$

As (4) is completely separable over i , we can optimise each summand independently:

$$\min_{L_i} \mathcal{L}_i(L_i) := \mathbb{E}_{q_i(\theta_i; L_i)}[-\log p(D_i|\theta_i)] + \mathbb{E}_{q(\phi; L_0)}[\text{KL}(q_i(\theta_i; L_i)||p(\theta_i|\phi))]. \quad (5)$$

So (5) constitutes local update/optimisation for client i . Note that each client i needs to access its private data D_i only, and not others, thus fully compatible with FL.

- Optimisation over L_0 (L_1, \dots, L_N fixed).

$$\min_{L_0} \mathcal{L}_0(L_0) := \text{KL}(q(\phi; L_0)||p(\phi)) - \sum_{i=1}^N \mathbb{E}_{q(\phi; L_0)q_i(\theta_i; L_i)}[\log p(\theta_i|\phi)]. \quad (6)$$

This constitutes server update criteria while the latest $q_i(\theta_i; L_i)$'s from local clients being fixed. Remarkably, the server needs not access any local data at all, suitable for FL. This nice property originates from the independence assumption in our approximate posterior (2).

Algorithm 1 Training algorithm: General framework.

Input: Initial parameters L_0 in the variational posterior $q(\phi; L_0)$.

Output: Trained parameters L_0 .

For each round $r = 1, 2, \dots, R$ **do**:

1. Sample a subset \mathcal{I} of participating clients ($|\mathcal{I}| = N_f \leq N$).
2. Server sends L_0 to all clients $i \in \mathcal{I}$.
3. **For** each client $i \in \mathcal{I}$ **in parallel do**:
 Solve (by SGD) with L_0 fixed:

$$\min_{L_i} \mathbb{E}_{q_i(\theta_i; L_i)}[-\log p(D_i|\theta_i)] + \mathbb{E}_{q(\phi; L_0)}[\text{KL}(q_i(\theta_i; L_i)||p(\theta_i|\phi))],$$

Initial L_i is either copied from L_0 or the last iterate if the client can save L_i locally.

4. Each client $i \in \mathcal{I}$ sends the updated L_i back to the server.
5. Upon receiving $\{L_i\}_{i \in \mathcal{I}}$, the server updates L_0 by solving (with $\{L_i\}_{i \in \mathcal{I}}$ fixed):

$$\min_{L_0} \text{KL}(q(\phi; L_0)||p(\phi)) - \frac{N}{N_f} \sum_{i \in \mathcal{I}} \mathbb{E}_{q(\phi; L_0)q_i(\theta_i; L_i)}[\log p(\theta_i|\phi)].$$

Pseudocodes. Our training algorithm (general framework) is summarised as pseudocodes in Alg. 1.

Interpretation. First, server’s loss function (6) tells us that the server needs to update $q(\phi; L_0)$ in such a way that (i) it puts mass on those ϕ that have high compatibility scores $\log p(\theta_i|\phi)$ with the current local models $\theta_i \sim q_i(\theta_i)$, thus aiming to be aligned with local models, and (ii) it does not deviate from the prior $p(\phi)$. Clients’ loss function (5) indicates that each client i needs to minimise the class prediction error on its own data D_i (first term), and at the same time, to stay close to the current global standard $\phi \sim q(\phi)$ by reducing the KL divergence from $p(\theta_i|\phi)$ (second term).

3.2 Formalisation of Global Prediction and Personalisation Tasks

Two important tasks in FL are: *global prediction* and *personalisation*. The former evaluates the trained model on novel test data sampled from a distribution possibly different from training data. Personalisation is the task of adapting the trained model on a new dataset called personalised data. In our Bayesian model, these two tasks can be formally defined as Bayesian inference problems.

Global prediction. The task is to predict the class label of a novel test input x^* which may or may not come from the same distributions as the training data D_1, \dots, D_N . Under our Bayesian model, it can be turned into a probabilistic inference problem $p(y^*|x^*, D_{1:N})$. Let θ be the local model that generates the output y^* given x^* . Exploiting conditional independence from Fig. 1(c),

$$p(y^*|x^*, D_{1:N}) = \iint p(y^*|x^*, \theta) p(\theta|\phi) p(\phi|D_{1:N}) d\theta d\phi \tag{7}$$

$$\approx \iint p(y^*|x^*, \theta) p(\theta|\phi) q(\phi) d\theta d\phi = \int p(y^*|x^*, \theta) \left(\int p(\theta|\phi) q(\phi) d\phi \right) d\theta, \tag{8}$$

Algorithm 2 Global prediction: General framework.

Input: Test input x^* . Learned model L_0 in the variational posterior $q(\phi; L_0)$.

Output: Predictive distribution $p(y^*|x^*, D_{1:N})$.

1. Sample $\theta^{(s)} \sim \int p(\theta|\phi) q(\phi; L_0) d\phi$ for $s = 1, \dots, S$.

2. Return $p(y^*|x^*, D_{1:N}) \approx \frac{1}{S} \sum_{s=1}^S p(y^*|x^*, \theta^{(s)})$.

where in (8) we use $p(\phi|D_{1:N}) \approx q(\phi)$. The inner integral (in parentheses) in (8) either admits a closed form (Sec. 4.1) or can be approximated (e.g., Monte-Carlo estimation). The pseudocode for the global prediction (general framework) is depicted in Alg. 2.

Personalisation. It formally refers to the task of learning a prediction model $\hat{p}(y|x)$ given an unseen (personal) training dataset D^p that comes from some unknown distribution $p^p(x, y)$, so that the personalised model \hat{p} performs well on novel (in-distribution) test points $(x^p, y^p) \sim p^p(x, y)$. Evidently we need to exploit (and benefit from) the trained model from the FL training stage. To this end many existing approaches simply resort to *finetuning*, that is, training on D^p warm-starting with the FL-trained model. However, a potential issue is the lack of a solid principle on how to balance the initial FL-trained model and personal data fitting to avoid underfitting and overfitting. In our Bayesian framework, the personalisation can be seen as another posterior inference problem with *additional evidence* of the personal training data D^p . Prediction on a test point x^p amounts to inferring:

$$p(y^p|x^p, D^p, D_{1:N}) = \int p(y^p|x^p, \theta) p(\theta|D^p, D_{1:N}) d\theta. \quad (9)$$

So, it boils down to the task of posterior inference $p(\theta|D^p, D_{1:N})$ given both the personal data D^p and the FL training data $D_{1:N}$. Under our hierarchical model, by exploiting conditional independence from graphical model (Fig. 1(d)), we can link the posterior to our FL-trained $q(\phi)$ as follows:

$$p(\theta|D^p, D_{1:N}) = \int p(\theta|D^p, \phi) p(\phi|D^p, D_{1:N}) d\phi \approx \int p(\theta|D^p, \phi) p(\phi|D_{1:N}) d\phi \quad (10)$$

$$\approx \int p(\theta|D^p, \phi) q(\phi) d\phi \quad (11)$$

$$\approx p(\theta|D^p, \phi^*). \quad (12)$$

In (10) we disregard the impact of D^p on the higher-level ϕ given the joint evidence, $p(\phi|D^p, D_{1:N}) \approx p(\phi|D_{1:N})$ due to the dominance of $D_{1:N}$ compared to smaller D^p . In (11) our variational posterior approximation is used, that is, $p(\phi|D_{1:N}) \approx q(\phi)$. Lastly, (12) makes approximation using the mode ϕ^* of $q(\phi)$, which is reasonable for our two modeling choices for $q(\phi)$ discussed in Sec. 4.1 (Sec. A.2) and Sec. 4.2 (Sec. A.3). Since dealing with $p(\theta|D^p, \phi^*)$ involves difficult marginalisation $p(D^p|\phi^*) = \int p(D^p|\theta)p(\theta|\phi^*)d\theta$, we adopt variational inference, introducing a tractable variational distribution $v(\theta) \approx p(\theta|D^p, \phi^*)$. Following the usual variational inference derivations, we have the negative ELBO objective (for personalisation):

$$\min_v \mathbb{E}_{v(\theta)}[-\log p(D^p|\theta)] + \text{KL}(v(\theta)||p(\theta|\phi^*)). \quad (13)$$

Algorithm 3 Personalisation: General framework.

Input: Personal training data D^p . Test input x^p . Learned model L_0 in the variational posterior $q(\phi; L_0)$.

Output: Predictive distribution $p(y^p|x^p, D^p, D_{1:N})$.

1. Estimate the variational density $v(\theta) \approx p(\theta|D^p, \phi^*)$ by solving (via SGD):

$$\min_v \mathbb{E}_{v(\theta)}[-\log p(D^p|\theta)] + \text{KL}(v(\theta)||p(\theta|\phi^*)), \text{ where } \phi^* = \arg \max_{\phi} q(\phi; L_0).$$

2. Sample $\theta^{(s)} \sim v(\theta)$ for $s = 1, \dots, S$.

3. Return $p(y^p|x^p, D^p, D_{1:N}) \approx \frac{1}{S} \sum_{s=1}^S p(y^p|x^p, \theta^{(s)})$.
-

With the optimised v , our predictive distribution becomes (S = the number of MC samples):

$$p(y^p|x^p, D^p, D_{1:N}) \approx \frac{1}{S} \sum_{s=1}^S p(y^p|x^p, \theta^{(s)}), \text{ where } \theta^{(s)} \sim v(\theta), \quad (14)$$

which simply requires feed-forwarding test input x^p through the sampled networks $\theta^{(s)}$ and averaging. The pseudocode for the personalisation (general framework) is depicted in Alg. 3.

Thus far, we have discussed a general framework, deriving how the variational inference for our Bayesian model fits gracefully in the FL problem. In the next section, we define specific density families for the prior $(p(\phi), p(\theta_i|\phi))$ and posterior $(q(\phi), q_i(\theta_i))$ as our proposed concrete models.

4. Bayesian FL: Two Concrete Models

We propose two different model choices that we find the most interesting: **Normal-Inverse-Wishart** (Sec. 4.1) and **Mixture** (Sec. 4.2). To avoid distraction, we make this section concise putting only the final results and discussions, and leaving all mathematical details in Appendix A.2 and A.3.

To recapitulate the notations in our hierarchical Bayesian modeling: we have two types of latent random variables, ϕ and $\{\theta_i\}_{i=1}^N$ where each θ_i is deployed as the network weights for client i 's backbone, and ϕ is a globally shared variable responsible for linking the individual client parameters θ_i s. The prior is decomposed as: $p(\phi, \{\theta_i\}_{i=1}^N) = p(\phi) \prod_{i=1}^N p(\theta_i|\phi)$. The posterior $p(\phi, \theta_{1:N}|D_{1:N})$ is approximated by the variational distribution, $q(\phi, \theta_{1:N}; L) = q(\phi; L_0) \prod_{i=1}^N q_i(\theta_i; L_i)$, where we optimize the variational parameters $L = \{L_0, \{L_i\}_{i=1}^N\}$ by minimising the negative ELBO loss via our block-coordinate-descent (FL) algorithm. The following two sections instantiate specifically these prior and variational distributions as well as the resulting client/server loss functions.

4.1 Normal-Inverse-Wishart (NIW) Model

We define the prior as a conjugate form of Gaussian and Normal-Inverse-Wishart. Specifically we let $\phi = (\mu, \Sigma)$ and,

$$p(\phi) = \mathcal{NIW}(\mu, \Sigma; \Lambda) = \mathcal{N}(\mu; \mu_0, \lambda_0^{-1}\Sigma) \cdot \mathcal{IW}(\Sigma; \Sigma_0, \nu_0), \quad (15)$$

$$p(\theta_i|\phi) = \mathcal{N}(\theta_i; \mu, \Sigma), \quad i = 1, \dots, N, \quad (16)$$

where $\Lambda = \{\mu_0, \Sigma_0, \lambda_0, \nu_0\}$ is the parameters of the NIW. Although Λ can be learned via data marginal likelihood maximisation (e.g., empirical Bayes), but for simplicity we leave it fixed as²: $\mu_0 = 0$, $\Sigma_0 = I$, $\lambda_0 = 1$, and $\nu_0 = d + 2$ where d is the number of parameters in θ_i or μ . Next, our choice of the variational density family for $q(\phi)$ is the NIW, not just because it is the most popular parametric family for a pair of mean vector and covariance matrix $\phi = (\mu, \Sigma)$, but it can also admit closed-form expressions in the ELBO function due to the conjugacy as we derive in Appendix A.2.1.

$$q(\phi) := \mathcal{N}\mathcal{I}\mathcal{W}(\phi; \{m_0, V_0, l_0, n_0\}) = \mathcal{N}(\mu; m_0, l_0^{-1}\Sigma) \cdot \mathcal{I}\mathcal{W}(\Sigma; V_0, n_0). \quad (17)$$

Although the scalar parameters l_0, n_0 can be optimised together with m_0, V_0 , their impact is less influential and we find that they make the ELBO optimisation a little bit cumbersome. So we fix l_0, n_0 with some near-optimal values by exploiting the conjugacy of the NIW under Gaussian likelihood (details in Appendix A.2), and regard m_0, V_0 as variational parameters, $L_0 = \{m_0, V_0\}$. We restrict V_0 to be diagonal for computational tractability. The density family for $q_i(\theta_i)$'s can be a Gaussian, but we find that it is computationally more attractive and numerically more stable to adopt the mixture of two spiky Gaussians that leads to the MC-Dropout (Gal and Ghahramani, 2016). That is,

$$q_i(\theta_i) = \prod_l (p_{do} \cdot \mathcal{N}(\theta_i[l]; m_i[l], \epsilon^2 I) + (1 - p_{do}) \cdot \mathcal{N}(\theta_i[l]; 0, \epsilon^2 I)), \quad (18)$$

where (i) m_i is the only variational parameters ($L_i = \{m_i\}$), (ii) $\cdot[l]$ indicates a column/layer in neural network parameters where l goes over layers and columns of weight matrices, (iii) p_{do} is the (user-specified) hyperparameter where $1 - p_{do}$ corresponds to the dropout probability, and (iv) ϵ is small constant (e.g., 10^{-4}) that makes two Gaussians spiky, close to the delta functions.

Client update. We apply the general client update optimisation (5) to the NIW model. Following the approximation of (Gal and Ghahramani, 2016) for the KL divergence between a mixture of Gaussians (18) and a Gaussian (16), we have the client local optimisation (details in Appendix A.2):

$$\min_{m_i} \mathcal{L}_i(m_i) := -\log p(D_i | \tilde{m}_i) + \frac{p_{do}}{2} (n_0 + d + 1) (m_i - m_0)^\top V_0^{-1} (m_i - m_0), \quad (19)$$

where \tilde{m}_i is the dropout version of m_i , i.e., a reparametrised sample from (18). Note that m_0 and V_0 are fixed during the optimisation. Interestingly (19) generalises Fed-Avg (McMahan et al., 2017) and Fed-Prox (Li et al., 2018): With $p_{do} = 1$ (i.e., no dropout) and setting $V_0 = \alpha I$, (19) reduces to the client update formula for Fed-Prox where constant α controls the impact of the proximal term.

Server update. The general server optimisation (6) admits the closed-form solution (Appendix A.2):

$$m_0^* = \frac{p_{do}}{N+1} \sum_i m_i, \quad V_0^* = \frac{n_0}{N+d+2} \left((1+N\epsilon^2)I + m_0^*(m_0^*)^\top + \sum_i \rho(m_0^*, m_i, p_{do}) \right), \quad (20)$$

where $\rho(m_0, m_i, p_{do}) = p_{do}m_i m_i^\top - p_{do}m_0 m_i^\top - p_{do}m_i m_0^\top + m_0 m_0^\top$. Note that m_i 's are fixed from clients' latest variational parameters. It is interesting to see that m_0^* in (20) generalises

2. This choice ensures that the mean of Σ equals I , and μ is distributed as 0-mean, Σ -covariance Gaussian.

Algorithm 4 Training algorithm: Normal-Inverse-Wishart case.

Input: Initial $L_0 = (m_0, V_0)$ in $q(\phi; L_0) = \mathcal{N}\mathcal{IW}(\phi; \{m_0, V_0, l_0 = |D|+1, n_0 = |D|+d+2\})$ where $|D| = \sum_{i=1}^N |D_i|$ and d is the number of parameters in the backbone network $p(y|x, \theta)$.

Output: Trained parameters $L_0 = (m_0, V_0)$.

For each round $r = 1, 2, \dots, R$ **do**:

1. Sample a subset \mathcal{I} of participating clients ($|\mathcal{I}| = N_f \leq N$).
2. Server sends $L_0 = (m_0, V_0)$ to all clients $i \in \mathcal{I}$.
3. **For** each client $i \in \mathcal{I}$ **in parallel do**:
Solve (by SGD) with $L_0 = (m_0, V_0)$ fixed:

$$\min_{m_i} -\log p(D_i|\tilde{m}_i) + \frac{p_{do}}{2}(n_0 + d + 1)(m_i - m_0)^\top V_0^{-1}(m_i - m_0),$$

where \tilde{m}_i is the dropout version (with probability $1 - p_{do}$) of m_i .

Initial m_i can be either copied from m_0 or the last iterate if the client can save m_i locally.

4. Each client $i \in \mathcal{I}$ sends the updated $L_i = m_i$ back to the server.
5. Upon receiving $\{m_i\}_{i \in \mathcal{I}}$, the server updates $L_0 = (m_0, V_0)$ by:

$$m_0^* = \frac{p_{do}}{N+1} \frac{N}{N_f} \sum_{i \in \mathcal{I}} m_i, \quad V_0^* = \frac{n_0}{N+d+2} \left((1 + N\epsilon^2)I + m_0^*(m_0^*)^\top + \frac{N}{N_f} \sum_{i \in \mathcal{I}} \rho(m_0^*, m_i, p_{do}) \right),$$

$$\text{where } \rho(m_0, \bar{m}_i, p_{do}) = p_{do}m_i m_i^\top - p_{do}m_0 m_i^\top - p_{do}m_i m_0^\top + m_0 m_0^\top.$$

Algorithm 5 Global prediction: Normal-Inverse-Wishart case.

Input: Test input x^* . Learned model $L_0 = (m_0, V_0)$ in $q(\phi; L_0) = \mathcal{N}\mathcal{IW}(\phi; \{m_0, V_0, l_0, n_0\})$.

Output: Predictive distribution $p(y^*|x^*, D_{1:N})$.

1. Sample $\theta^{(s)} \sim t_{n_0-d+1} \left(\theta; m_0, \frac{(l_0+1)V_0}{l_0(n_0-d+1)} \right)$ for $s = 1, \dots, S$.
 2. Return $p(y^*|x^*, D_{1:N}) \approx \frac{1}{S} \sum_{s=1}^S p(y^*|x^*, \theta^{(s)})$.
-

the well-known aggregation step of averaging local models in Fed-Avg (McMahan et al., 2017) and related methods: when $p_{do}=1$ (no dropout), it almost equals client model averaging. Also, since $\rho(m_0^*, m_i, p_{do}=1) = (m_i - m_0^*)(m_i - m_0^*)^\top$ when $p_{do} = 1$, V_0^* essentially estimates the sample scatter matrix with $(N+1)$ samples, namely clients' m_i 's and server's prior $\mu_0 = 0$, measuring how much they deviate from the center m_0^* . The dropout is known to help regularise the model and lead to better generalisation (Gal and Ghahramani, 2016), and with $p_{do} < 1$ our (20) forms a principled optimal solution.

The training algorithm for our NIW model is summarised as pseudocodes in Alg. 4.

Global prediction. The inner integral of (8) becomes the multivariate Student- t distribution. Then the predictive distribution for a new test input x^* can be estimated as (usually $S = 1$ in practice):

$$p(y^*|x^*, D_{1:N}) \approx \frac{1}{S} \sum_{s=1}^S p(y^*|x^*, \theta^{(s)}), \quad \theta^{(s)} \sim t_{n_0-d+1} \left(\theta; m_0, \frac{(l_0+1)V_0}{l_0(n_0-d+1)} \right), \quad (21)$$

where $t_\nu(a, B)$ is the multivariate Student- t with location a , scale matrix B , and d.o.f. ν . The pseudocode for the global prediction in the NIW model is depicted in Alg. 5.

Algorithm 6 Personalisation: Normal-Inverse-Wishart case.

Input: Personal training data D^p . Test input x^p .

Learned model $L_0 = (m_0, V_0)$ in $q(\phi; L_0) = \mathcal{N}\mathcal{W}(\phi; \{m_0, V_0, l_0, n_0\})$.

Output: Predictive distribution $p(y^p|x^p, D^p, D_{1:N})$.

1. Estimate m in $v(\theta; m) = \prod_l (p_{do} \cdot \mathcal{N}(\theta[l]; m[l], \epsilon^2 I) + (1-p_{do}) \cdot \mathcal{N}(\theta[l]; 0, \epsilon^2 I))$ by solving (SGD):

$$\min_m -\log p(D^p|\tilde{m}) + \frac{p_{do}}{2} (n_0 + d + 1)(m - m_0)^\top V_0^{-1}(m - m_0),$$

where \tilde{m} is the dropout version (with probability $1 - p_{do}$) of m .

2. Return $p(y^p|x^p, D^p, D_{1:N}) \approx p(y^p|x^p, m)$.
-

Personalisation. With the given personalisation training data D^p , we follow the general framework in (13) to find $v(\theta) \approx p(\theta|D^p, \phi^*)$ in a variational way, where ϕ^* obtained from (38). We adopt the same spiky mixture form (18) for $v(\theta)$, which leads to the learning objective similar to (19). The pseudocode for the personalisation in the NIW model is depicted in Alg. 6.

4.2 Mixture Model

Our motivation for mixture is to make the prior $p(\theta, \phi)$ more flexible by having multiple different prototypes, diverse enough to cover the heterogeneity in data distributions across clients. Specifically we let $\phi = \{\mu_1, \dots, \mu_K\}$, and define the prior as follows so that it contains K networks (prototypes) that can broadly cover the clients data distributions:

$$p(\phi) = \prod_{j=1}^K \mathcal{N}(\mu_j; 0, I), \quad p(\theta_i|\phi) = \sum_{j=1}^K \frac{1}{K} \mathcal{N}(\theta_i; \mu_j; \sigma^2 I), \quad (22)$$

where σ is the hyperparameter that captures perturbation scale, chosen by users. Note that we put equal mixing proportions $1/K$ due to the symmetry, a priori. That is, each client can take any of μ_j 's equally likely a priori. For the variational densities, we define:

$$q_i(\theta_i) = \mathcal{N}(\theta_i; m_i, \epsilon^2 I), \quad q(\phi) = \prod_{j=1}^K \mathcal{N}(\mu_j; r_j, \epsilon^2 I), \quad (23)$$

where $\{r_j\}_{j=1}^K$ (L_0), m_i (L_i) are variational parameters, and ϵ is small constant (e.g., 10^{-4}).

Client update. The general client update (5) reduces to (details in Appendix A.3):

$$\min_{m_i} \mathbb{E}_{q_i(\theta_i)} [-\log p(D_i|\theta_i)] - \log \sum_{j=1}^K \exp\left(-\frac{\|m_i - r_j\|^2}{2\sigma^2}\right). \quad (24)$$

It is interesting to see that (24) can be seen as generalisation of Fed-Prox (Li et al., 2018), where the proximal regularisation term in Fed-Prox is extended to *multiple* global models r_j 's, penalizing the local model (m_i) straying away from these prototypes. And if we use a single prototype ($K = 1$), the optimisation (24) exactly reduces to the local update objective of Fed-Prox. Since **log-sum-exp** is approximately equal to **max**, the regularisation term in (24) effectively focuses on the closest global prototype r_j from the current local model m_i , which is intuitively well aligned with our motivation.

Server update. The general form (6) can be approximately turned into (Appendix A.3 for derivations):

$$\min_{\{r_j\}_{j=1}^K} \frac{1}{2} \sum_{j=1}^K \|r_j\|^2 - \sum_{i=1}^N \log \sum_{j=1}^K \exp\left(-\frac{\|m_i - r_j\|^2}{2\sigma^2}\right). \quad (25)$$

Interestingly, (25) generalises the well-known aggregation step of averaging local models in Fed-Avg and related methods: Especially when $K = 1$, (25) reduces to quadratic optimisation, admitting the optimal solution $r_1^* = \frac{1}{N+\sigma^2} \sum_{i=1}^N m_i$. The extra term σ^2 can be explained by incorporating an extra *zero* local model originating from the prior (interpreted as a *neutral* model) with the discounted weight σ^2 rather than 1. Although (25) for $K > 1$ can be solved by standard gradient descent, we apply the Expectation-Maximisation (EM) algorithm³ (Dempster et al., 1977) instead:

$$\text{(E-step)} \quad c(j|i) = \frac{e^{-\|m_i - r_j\|^2/(2\sigma^2)}}{\sum_{j=1}^K e^{-\|m_i - r_j\|^2/(2\sigma^2)}}, \quad \text{(M-step)} \quad r_j^* = \frac{\frac{1}{N} \sum_{i=1}^N c(j|i) \cdot m_i}{\frac{\sigma^2}{N} + \frac{1}{N} \sum_{i=1}^N c(j|i)}. \quad (26)$$

The M-step (server update) has intuitive meaning that the new prototype r_j becomes the *weighted* average of the local models m_i 's where the weights $c(j|i)$ are determined by the proximity between m_i and r_j (i.e., those m_i 's that are closer to r_j have more contribution, and vice versa). This can be seen as an extension of the aggregation step in Fed-Avg to the multiple prototype case. The training algorithm for our mixture model is summarised as pseudocodes in Alg. 7.

Global prediction. We slightly modify our general approach to make client data dominantly explained by the most relevant model r_j , by introducing a gating network (Appendix A.3 for details). The pseudocode for the global prediction is depicted in Alg. 8.

Personalisation. With $v(\theta)$ of the same form as $q_i(\theta_i)$, the VI learning becomes similar to (24). The pseudocode for the personalisation in the mixture model is depicted in Alg. 9.

5. Theoretical Analysis

5.1 Convergence Analysis

As a special block-coordinate optimisation algorithm, we show that our FL algorithm converges to an (local) optimum of the training objective (3).

Theorem 1 (Convergence analysis). *We denote the objective function in (3) by $f(x)$ where $x = [x_0, x_1, \dots, x_N]$ corresponding to the variational parameters $x_0 := L_0$, $x_1 := L_1$, \dots , $x_N := L_N$. Let $\eta_t = \bar{L} + \sqrt{t}$ for some constant \bar{L} , and $\bar{x}^T = \frac{1}{T} \sum_{t=1}^T x^t$, where t is the batch iteration counter, x^t is the iterate at t by following our FL algorithm, and $N_f (\leq N)$ is the number of participating clients at each round. With Assumptions 1–3 in Appendix B.1, the following holds for any T :*

$$\mathbb{E}[f(\bar{x}^T)] - f(x^*) \leq \frac{N + N_f}{N_f} \cdot \frac{\frac{\sqrt{T+\bar{L}}}{2} D^2 + R_f^2 \sqrt{T}}{T} = O\left(\frac{1}{\sqrt{T}}\right), \quad (27)$$

3. Instead of performing several EM steps until convergence, in practice we find only one EM step is sufficient.

Algorithm 7 Training algorithm: Mixture case.

Input: Initial $L_0 = \{r_j\}_{j=1}^K$ in $q(\phi; L_0) = \prod_j \mathcal{N}(\mu_j; r_j, \epsilon^2 I)$ and β in the gating network $g(x; \beta)$.

Output: Trained parameters $L_0 = \{r_j\}_{j=1}^K$ and β .

For each round $r = 1, 2, \dots, R$ **do**:

1. Sample a subset \mathcal{I} of participating clients ($|\mathcal{I}| = N_f \leq N$).
2. Server sends $L_0 = \{r_j\}_{j=1}^K$ and β to all clients $i \in \mathcal{I}$.
3. **For** each client $i \in \mathcal{I}$ **in parallel do**:
 Solve (by SGD) with $L_0 = \{r_j\}_{j=1}^K$ fixed:

$$\min_{m_i} \mathbb{E}_{q_i(\theta_i; m_i)} [-\log p(D_i | \theta_i)] - \log \sum_j \exp \left(-\frac{\|m_i - r_j\|^2}{2\sigma^2} \right), \text{ where } q_i(\theta_i; m_i) = \mathcal{N}(\theta_i; m_i, \epsilon^2 I).$$

Initial m_i is either the center of $\{r_j\}$ or the last iterate if the client can save m_i locally.

$\beta_i = \text{SGD update of } \beta \text{ in } g(x; \beta) \text{ with data } \{(x, j^*)\}_{x \sim D_i}$ where $j^* = \arg \min_j \|m_i - r_j\|$.

4. Each client $i \in \mathcal{I}$ sends the updated $L_i = m_i$ and β_i back to the server.
5. Upon receiving $\{m_i\}_{i \in \mathcal{I}}$ and $\{\beta_i\}_{i \in \mathcal{I}}$, the server updates $L_0 = \{r_j\}_{j=1}^K$ by the one-step EM:

$$\text{(E-step)} \quad c(j|i) = \frac{e^{-\|m_i - r_j\|^2 / (2\sigma^2)}}{\sum_{j=1}^K e^{-\|m_i - r_j\|^2 / (2\sigma^2)}}, \quad \text{(M-step)} \quad r_j^* = \frac{\frac{1}{N_f} \sum_{i \in \mathcal{I}} c(j|i) \cdot m_i}{\frac{\sigma^2}{N} + \frac{1}{N_f} \sum_{i \in \mathcal{I}} c(j|i)},$$

and updates β by aggregation: $\beta^* = \frac{1}{N_f} \sum_{i \in \mathcal{I}} \beta_i$.

Algorithm 8 Global prediction: Mixture case.

Input: Test input x^* . Learned $L_0 = \{r_j\}_{j=1}^K$ in $q(\phi; L_0) = \prod_j \mathcal{N}(\mu_j; r_j, \epsilon^2 I)$ and β in $g(x; \beta)$.

Output: Predictive distribution $p(y^* | x^*, D_{1:N})$.

1. Return $p(y^* | x^*, D_{1:N}) \approx \sum_{j=1}^K g_j(x^*) \cdot p(y^* | x^*, r_j)$.
-

Algorithm 9 Personalisation: Mixture case.

Input: Personal training data D^p . Test input x^p . Learned model $L_0 = \{r_j\}_{j=1}^K$.

Output: Predictive distribution $p(y^p | x^p, D^p, D_{1:N})$.

1. Estimate m in $v(\theta; m) = \mathcal{N}(\theta; m, \epsilon^2 I)$ by solving (via SGD):

$$\min_m \mathbb{E}_{v(\theta; m)} [-\log p(D^p | \theta)] - \log \sum_j \exp \left(-\frac{\|m - r_j\|^2}{2\sigma^2} \right).$$

2. Return $p(y^p | x^p, D^p, D_{1:N}) \approx p(y^p | x^p, m)$.
-

where x^* is the (local) optimum, D , and R_f are some constants, and the expectation is taken over randomness in minibatches and selection of participating clients.

Proof See Appendix B.1. ■

Remark. Theorem 1 states that \bar{x}^t converges to the optimal point x^* in expectation at the rate of $O(1/\sqrt{t})$. This rate asymptotically equals that of the conventional (non-block-coordinate, holistic) SGD algorithm.

5.2 Generalisation Error Bound

We theoretically show how well this optimal model trained on empirical data performs on unseen test data points.

Theorem 2 (Generalisation error bound). *Assume that the variational density family for $q_i(\theta_i)$ is rich enough to subsume Gaussian. Let $d^2(P_{\theta_i}, P^i)$ be the expected squared Hellinger distance⁴ between the true class distribution $P^i(y|x)$ and model’s $P_{\theta_i}(y|x)$ for client i ’s data. The optimal solution $(\{q_i^*(\theta_i)\}_{i=1}^N, q^*(\phi))$ of the optimisation problem (3) satisfies:*

$$\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{q_i^*(\theta_i)}[d^2(P_{\theta_i}, P^i)] \leq O\left(\frac{1}{n}\right) + C \cdot \epsilon_n^2 + C' \left(r_n + \frac{1}{N} \sum_{i=1}^N \lambda_i^* \right), \quad (28)$$

with high probability, where $C, C' > 0$ are constant, $\lambda_i^* = \min_{\theta \in \Theta} \|f_{\theta} - f^i\|_{\infty}^2$ is the best error within our backbone network family Θ , and $r_n, \epsilon_n \rightarrow 0$ as the training data size $n \rightarrow \infty$.

Proof See Appendix B.2. ■

Remark. *Theorem 2 implies that the optimal solution of (3) (attainable by our block-coordinate FL algorithm) is asymptotically optimal, since the RHS of (28) converges to 0 as the training data size $n \rightarrow \infty$. Note that the last term $\frac{1}{N} \sum_i \lambda_i^*$ can be made arbitrarily close to 0 by increasing the backbone capacity (MLPs as universal function approximators). But practically for fixed n , as enlarging the backbone capacity (i.e., large T , L , and M) also increases ϵ_n and r_n (definitions of these sequences and details in Appendix B.2), it is important to choose the backbone network architecture properly. Note also that our assumption on the variational density family for $q_i(\theta_i)$ is easily met; for instance, the families of the mixtures of Gaussians adopted in NIW (Sec. 4.1) and mixture models (Sec. 4.2) obviously subsume a single Gaussian family.*

5.3 Computational Complexity Analysis

The computational complexity of the proposed algorithms is summarised in Table 1 (training complexity with communication costs), Table 2 (global prediction complexity), and Table 3 (personalisation complexity). Our methods incur only constant-factor extra cost compared to the minimal-cost FedAvg.

6. Evaluation

We evaluate the proposed hierarchical Bayesian models on several FL benchmarks: **CIFAR-100**, **MNIST**, **Fashion-MNIST**, and **EMNIST**. We also have the results on the challenging corrupted CIFAR (**CIFAR-C-100**) that renders the client data more heterogeneous both in input images and class distributions. Our implementation⁵ is based on (Oh et al.,

4. It is defined as $d^2(P_{\theta}, P^i) = \mathbb{E}_{x \sim P^i(x)} [H^2(P_{\theta}(y|x), P^i(y|x))] = \mathbb{E}_{x \sim P^i(x)} \left[1 - \exp\left(-\frac{\|f_{\theta}(x) - f^i(x)\|_2^2}{8\sigma_{\epsilon}^2}\right) \right]$.

5. The codes to reproduce the results can be found in the Supplement.

Table 1: **Training complexity** of the proposed algorithms (NIW and Mixture) and FedAvg. All quantities are per-round, per-batch, and per-client costs. In the entries, d = the number of parameters in the backbone network, F = time for feed-forward pass, B = time for backprop, and N_f = the number of participating clients per round.

	Communication cost		Client Local update	Server update
	Server \rightarrow Client	Client \rightarrow Server		
NIW	$2d$ (sent: m_0, V_0)	d (sent: m_i)	$F + B + O(d)$ ($O(d)$ from quadratic penalty)	$O(N_f \cdot d)$
Mixture (Order K)	$(K + 1)d$ (sent: $\{r_j\}_{j=1}^K, \beta$)	$2d$ (sent: m_i, β_i)	$2(F + B) + O(K \cdot d)$ ($O(K \cdot d)$ from log-sum-exp)	$O(K \cdot N_f \cdot d)$
FedAvg	d (sent: θ)	d (sent: θ_i)	$F + B$	$O(N_f \cdot d)$ (aggregation)

Table 2: **Global prediction complexity** of the proposed algorithms (NIW and Mixture) and FedAvg. All quantities are per-test-batch costs. In the entries, d = the number of parameters in the backbone network, F = time for feed-forward pass, and S = the number of samples $\theta^{(s)}$ from the Student- t distribution in the NIW case (we use $S = 1$).

Per-test-batch complexity	
NIW	$S \cdot F + O(S \cdot d)$ ($O(S \cdot d)$ from the cost of t -sampling)
Mixture (Order K)	$(K + 1)F$ (a forward pass for the gating network)
FedAvg	F

Table 3: **Personalisation complexity** of the proposed algorithms (NIW and Mixture) and FedAvg. All quantities are per-train/test-batch costs. In the entries, d = the number of parameters in the backbone network, F = time for feed-forward pass, and B = time for backprop.

	Training (personalisation) complexity	Test complexity
NIW	$F + B + O(d)$ ($O(d)$ from quadratic penalty)	F
Mixture (Order K)	$F + B + O(K \cdot d)$ ($O(K \cdot d)$ from log-sum-exp)	F
FedAvg	$F + B$	F

2022) where MobileNet (Howard et al., 2017) is used as a backbone, and follow the body-update strategy: the classification head (the last layer) is randomly initialised and fixed during training, with only the network body updated (and both body and head updated during personalisation). We report results all based on this body-update strategy since we

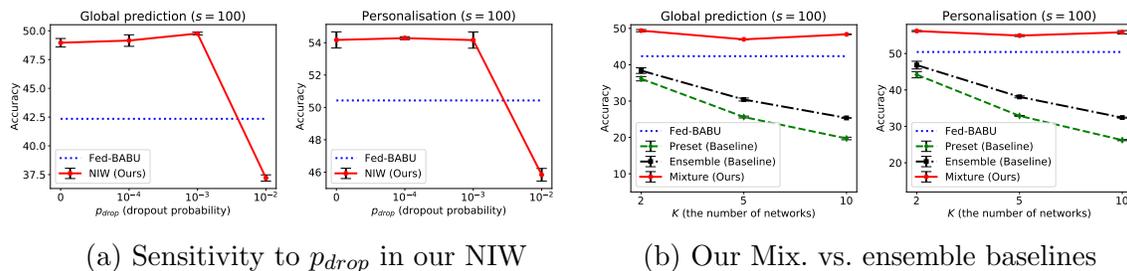


Figure 2: Hyperparameter sensitivity analysis and comparison with simple ensemble baselines.

observe that it considerably outperforms the full update for our models and other competing methods. The hyperparameters are: (**NIW**) $\epsilon = 10^{-4}$ and $p_{do} = 1 - 0.001$ (See Sec. 6.4.3 for other values); (**Mixture**) $\sigma^2 = 0.1$, $\epsilon = 10^{-4}$, mixture order $K=2$ (See Sec. 6.4.1 for other values), and the gating network has the same architecture as the main backbone, but the output cardinality changed to K . Other hyperparameters including batch size (50), learning rate (0.1 initially, decayed by 0.1) and the number of epochs in personalisation (5), are the same as those in (Oh et al., 2022).

Evaluation Protocols. Our experiment proceeds with the following three stages:

- **Training stage:** With a server and clients with their own local data, we perform our block-coordinate variational inference optimisation as described in Alg. 1 (Alg. 4 for NIW and Alg. 7 for Mixture). At the end of training, we maintain the learned variational distribution of the server variable $q(\phi; L_0)$, which is required in the subsequent global prediction and personalisation stages.
- **Global prediction stage:** The purpose of this stage is to test the performance of the learned global model $q(\phi; L_0)$ on test data points x^* . As described in Alg. 2 (Alg. 5 for NIW and Alg. 8 for Mixture), we build a predictive model using the learned $q(\phi; L_0)$ and the prior $p(\theta|\phi)$. In our experiments, we use the test data splits of the clients that participate in the training stage as probing test points.
- **Personalisation stage:** For a new user with their own local training data, typical personalisation amounts to fine-tuning the trained backbone with the user’s local training data. In our case, we have the Bayesian posterior inference driven personalisation algorithm as specified in Alg. 3 (Alg. 6 for NIW and Alg. 9 for Mixture), which runs another variational inference with the posterior $q(\phi; L_0)$ now serving as a prior. In our experiments novel users are created in a similar manner as the training stage, and once learned, the test data splits of the novel users are used for evaluating test performance.

Competing approaches are listed as follows:

- Fed-BABU (Oh et al., 2022)
- Fed-Avg (McMahan et al., 2017)
- Fed-Prox (Li et al., 2018)
- pFedBayes (Zhang et al., 2022)

- FedPA (Al-Shedivat et al., 2021)
- FedBE (Chen and Chao, 2021)
- FedEM (Marfoq et al., 2021)
- FedPop (Kotelevskii et al., 2022)

6.1 Datasets and Settings

6.1.1 CIFAR-100

Following (Oh et al., 2022), the client data distributions are heterogeneous non-iid, formed by the sharding-based class sampling (McMahan et al., 2017). More specifically, we partition data instances in each class into non-overlapping equal-sized shards, and assign s randomly sampled shards (over all classes) to each of N clients. Thus the number of shards per user s can control the degree of data heterogeneity: small s leads to more heterogeneity, and vice versa. The number of clients $N = 100$ (each having 500 training, 100 test samples), and we denote by f the fraction of participating clients. So, $N_f = \lfloor N \cdot f \rfloor$ clients are randomly sampled at each round to participate in training. Smaller f makes the FL more challenging, and we test two settings: $f = 1.0$ and 0.1 . Lastly, the number of epochs for client local update at each round is denoted by τ where we test $\tau = 1$ and 10 , and the number of total rounds is determined by τ as $\lfloor 320/\tau \rfloor$ for fairness. Note that smaller τ incurs more communication cost but often leads to higher accuracy. The competing FedBE (Chen and Chao, 2021) and FedEM (Marfoq et al., 2021), we set the number of ensemble components or base models to 3. FedPA (Al-Shedivat et al., 2021): shrinkage parameter $\rho = 0.01$.

6.1.2 CIFAR-100-CORRUPTED (CIFAR-C-100)

The CIFAR-100-Corrupted dataset (Hendrycks and Dietterich, 2019) makes CIFAR-100’s test split (10K images) corrupted by 19 different noise processes (e.g., Gaussian, motion blur, JPEG). For each corruption type, there are 5 corruption levels, and we use the severest one. Randomly chosen 10 corruption types are used for training (fixed) and the rest 9 types for personalisation. We divide $N = 100$ clients into 10 groups, each group assigned one of the 10 training corruption types exclusively (denoted by D^c the corrupted data for the group $c = 1, \dots, 10$). Each D^c is partitioned into 90%/10% training/test splits, and clients in each group ($\lfloor N/10 \rfloor$ clients) gets non-iid train/test subsets from D^c ’s train/test splits by following the sharding strategy with $s = 100$ or 50 . This way, the clients in different groups have considerable distribution shift in input images, while there also exists heterogeneity in class distributions even within the same groups.

For the FL-trained models, we evaluate *global prediction* on two datasets: clients’ test splits from the 10 training corruption types and the original (uncorrupted) CIFAR’s training split (50K images). For *personalisation*, we partition the clients into 9 groups, and assign one of the 9 corruption types to each group exclusively. Within each group we form non-iid sharding-based subsets similarly, and again we split the data into the 90% training/finetuning split and 10% test. Note that this personalisation setting is more challenging compared to CIFAR-100 since the data for personalisation are utterly unseen during the FL training stage. We test $\tau = 1$ and 4 scenarios. We test sharding parameter $s = 100$ or 50 , participating client fraction $f = 1.0$ or 0.1 , and the number of local epochs $\tau = 1$ and 4 scenarios.

Table 4: (CIFAR-100) Global prediction and personalisation test accuracy. We vary the FL setting (s, f, τ) where s controls the degree of data heterogeneity (small/large s for more/less heterogeneity), f is the fraction of participating clients, and τ is the number of epochs for client local update at each round.

(a) Global prediction performance (initial accuracy)												
FL setting			Our Methods		Fed-BABU	Fed-Avg	Fed-Prox	pFedBayes	FedPA	FedBE	FedEM	FedPop
s	f	τ	NIW	Mix ($K=2$)								
100	0.1	1	49.76±0.12	49.37±0.30	42.35±0.42	40.87±0.62	41.49±0.75	37.23±0.88	42.15±0.78	42.49±0.89	43.83±1.07	43.09±0.30
		10	29.02±0.33	29.02±0.29	27.93±0.28	28.26±0.19	27.11±0.11	28.21±1.42	28.05±0.28	28.39±0.45	28.62±0.26	28.31±0.42
	1.0	1	57.80±0.10	52.94±0.36	48.17±0.56	47.44±0.20	47.66±1.49	44.89±0.32	47.96±0.17	48.69±0.70	50.28±0.72	48.36±0.44
		10	29.53±0.42	30.55±0.15	28.67±0.51	28.79±0.68	27.43±0.38	28.25±0.81	28.89±0.38	28.60±1.18	29.51±0.12	28.99±0.47
10	0.1	1	37.54±0.25	38.07±0.40	35.04±0.56	27.48±0.86	34.73±0.21	31.49±0.18	35.51±0.55	35.17±0.40	37.28±0.26	35.01±0.58
		10	18.99±0.03	18.95±0.13	18.54±0.37	14.69±0.40	16.84±0.48	17.93±0.68	18.59±0.19	18.67±0.09	18.45±0.10	18.68±0.12
	1.0	1	50.40±0.11	49.52±0.88	45.41±0.11	37.10±0.44	44.33±0.31	39.95±0.89	45.08±0.72	45.56±0.52	47.52±0.59	44.98±0.27
		10	22.87±0.41	23.59±0.47	21.92±0.66	17.38±0.32	19.54±0.38	21.85±0.50	22.60±0.29	21.73±0.64	22.51±0.06	22.06±0.40

(b) Personalisation performance												
FL setting			Our Methods		Fed-BABU	Fed-Avg	Fed-Prox	pFedBayes	FedPA	FedBE	FedEM	FedPop
s	f	τ	NIW	Mix ($K=2$)								
100	0.1	1	54.16±0.50	56.17±0.16	50.43±0.93	46.43±0.82	49.91±0.78	45.83±1.12	49.88±0.49	50.57±1.03	47.28±0.88	51.22±0.37
		10	36.68±0.37	36.32±0.27	35.45±0.34	33.57±0.06	33.92±0.22	35.74±1.36	35.06±0.30	35.51±0.62	34.41±1.13	35.69±0.47
	1.0	1	60.36±0.89	58.82±0.37	55.87±0.91	53.15±0.25	55.50±0.90	53.00±0.48	55.43±0.31	56.25±0.42	54.65±0.49	55.85±0.51
		10	35.92±0.17	36.22±0.17	35.58±0.24	33.82±1.04	33.70±0.42	35.57±1.02	35.21±0.50	34.92±0.94	35.21±0.42	35.36±0.36
10	0.1	1	79.41±0.24	79.70±0.19	75.44±0.36	70.36±1.02	75.06±0.67	73.93±0.14	75.76±0.36	76.19±0.42	75.52±0.50	74.97±0.72
		10	67.35±1.02	67.57±0.62	66.24±0.53	61.39±0.27	64.86±0.73	65.82±0.33	65.87±0.21	66.64±0.25	67.11±0.11	66.70±0.55
	1.0	1	82.71±0.37	81.03±0.35	78.92±0.23	76.98±0.66	78.56±0.55	78.08±0.28	78.84±0.16	79.82±0.52	80.65±0.32	78.96±0.14
		10	67.78±1.02	66.74±0.27	66.25±0.46	63.81±0.40	63.81±0.51	66.15±1.29	66.23±0.29	66.06±0.54	66.34±0.12	66.57±0.44

6.1.3 MNIST, FASHION-MNIST AND EMNIST

Following the standard protocols, we set the number of clients $N = 100$, the number of shards per client $s = 5$, the fraction of participating clients per round $f = 0.1$, and the number of local training epochs per round $\tau = 1$ (total number of rounds 100) or 5 (total number of rounds 20) for MNIST and F-MNIST. For EMNIST, we have $N = 200$, $f = 0.2$, $\tau = 1$ (total number of rounds 300). We follow the standard Dirichlet-based client data splitting. For the competing methods FedBE (Chen and Chao, 2021) and FedEM (Marfoq et al., 2021), we use three-component models. The backbone is an MLP with a single hidden layer with 256 units for MNIST/F-MNIST, while we use a standard ConvNet with two hidden layers for EMNIST.

6.2 Main Results and Interpretation

In Table 4 (CIFAR-100), Table 5 and 6, we compare our methods (NIW and Mixture with $K=2$) against the popular FL methods, including FedAvg (McMahan et al., 2017), FedBABU (Oh et al., 2022), FedProx (Li et al., 2018), as well as recent Bayesian/ensemble methods, FedPA (Al-Shedivat et al., 2021), FedBE (Chen and Chao, 2021), pFedBayes (Zhang et al., 2022). FedEM (Marfoq et al., 2021), and FedPop (Kotelevskii et al., 2022) (See Sec. 2). We run the competing methods (implementation based on their public codes or our own

Table 5: (CIFAR-C-100) Global prediction and personalisation accuracy.

(a) Global prediction (initial accuracy) on test splits for the 10 training corruption types

FL settings			Our Methods		Fed-BABU	Fed-Avg	Fed-Prox	pFedBayes
s	f	τ	NIW	Mix. ($K = 2$)				
100	0.1	1	81.22\pm0.14	80.34 \pm 1.44	79.45 \pm 0.71	70.01 \pm 0.77	78.94 \pm 0.82	71.14 \pm 0.33
		4	67.69\pm0.74	65.81 \pm 0.84	63.58 \pm 1.28	48.47 \pm 1.26	60.96 \pm 1.11	57.88 \pm 1.51
	1.0	1	91.26\pm0.83	86.84 \pm 0.22	86.84 \pm 0.83	85.10 \pm 0.18	87.03 \pm 0.55	82.80 \pm 1.28
		4	73.58 \pm 1.02	74.55\pm0.22	71.03 \pm 0.75	58.32 \pm 0.07	65.76 \pm 0.10	64.07 \pm 1.89
50	0.1	1	78.63 \pm 0.39	79.36\pm0.24	77.44 \pm 1.17	68.27 \pm 0.53	78.31 \pm 0.93	67.95 \pm 0.22
		4	65.08\pm1.75	63.52 \pm 0.48	62.65 \pm 0.12	43.57 \pm 1.99	58.70 \pm 1.89	54.20 \pm 1.30
	1.0	1	89.31\pm0.17	88.24 \pm 0.71	86.44 \pm 0.97	83.31 \pm 0.67	86.00 \pm 0.93	82.32 \pm 0.37
		4	70.33\pm0.18	70.19 \pm 1.41	67.66 \pm 0.78	52.48 \pm 1.06	60.72 \pm 1.26	60.17 \pm 0.23

(b) Global prediction (initial accuracy) on the original (uncorrupted) CIFAR-100 training sets

FL settings			Our Methods		Fed-BABU	Fed-Avg	Fed-Prox	pFedBayes
s	f	τ	NIW	Mix. ($K = 2$)				
100	0.1	1	41.55\pm0.11	36.99 \pm 0.09	34.76 \pm 0.50	34.40 \pm 0.31	35.44 \pm 0.66	35.78 \pm 0.73
		4	30.84\pm0.07	30.60 \pm 0.27	28.31 \pm 0.28	29.24 \pm 0.79	28.09 \pm 0.71	29.12 \pm 0.30
	1.0	1	41.32\pm0.32	38.35 \pm 0.73	35.58 \pm 0.41	36.34 \pm 0.14	36.32 \pm 0.58	37.37 \pm 0.62
		4	30.67\pm0.12	30.40 \pm 0.44	28.60 \pm 0.28	30.31 \pm 0.87	27.95 \pm 0.28	29.14 \pm 0.57
50	0.1	1	41.04\pm0.14	36.41 \pm 0.47	35.44 \pm 0.58	34.13 \pm 0.50	36.37 \pm 0.50	35.68 \pm 0.27
		4	32.29\pm0.36	31.50 \pm 0.34	29.68 \pm 0.08	29.19 \pm 0.14	29.20 \pm 0.50	30.10 \pm 0.44
	1.0	1	41.64\pm0.21	38.54 \pm 0.42	36.09 \pm 0.28	35.78 \pm 0.83	37.13 \pm 0.55	38.39 \pm 0.21
		4	32.17\pm0.48	30.68 \pm 0.46	29.28 \pm 0.17	30.45 \pm 0.45	28.73 \pm 0.26	29.74 \pm 0.25

(c) Personalisation performance on the 9 held-out corruption types

FL settings			Our Methods		Fed-BABU	Fed-Avg	Fed-Prox	pFedBayes
s	f	τ	NIW	Mix. ($K = 2$)				
100	0.1	1	72.63 \pm 2.13	74.16\pm3.04	69.93 \pm 1.24	62.32 \pm 1.46	72.94 \pm 0.37	62.52 \pm 4.01
		4	62.74\pm0.94	61.56 \pm 1.69	60.33 \pm 2.12	53.35 \pm 1.39	56.91 \pm 1.40	53.71 \pm 2.13
	1.0	1	83.62 \pm 1.84	84.88\pm0.85	77.55 \pm 2.05	83.44 \pm 1.68	80.96 \pm 3.18	72.44 \pm 0.13
		4	64.84 \pm 1.05	67.35\pm1.46	53.25 \pm 2.19	53.34 \pm 1.22	41.39 \pm 1.18	43.41 \pm 1.92
50	0.1	1	75.50\pm0.74	67.33 \pm 2.83	59.47 \pm 2.48	57.77 \pm 0.80	56.34 \pm 2.80	53.47 \pm 0.51
		4	44.90 \pm 1.23	46.39\pm0.83	44.74 \pm 1.30	44.60 \pm 2.36	39.94 \pm 1.81	37.24 \pm 1.80
	1.0	1	81.46 \pm 0.67	81.77\pm3.11	67.43 \pm 2.58	72.52 \pm 1.02	71.52 \pm 2.02	68.47 \pm 1.43
		4	50.84\pm0.74	48.90 \pm 0.51	40.87 \pm 3.01	45.10 \pm 0.50	41.95 \pm 1.28	41.27 \pm 0.14

implementation if unavailable) with default hyperparameters (e.g., $\mu = 0.01$ for FedProx) and report the results. First of all, our two models (NIW and Mix.) consistently perform the best (by large margins most of the time) in terms of both global prediction and personalisation for nearly all FL settings on the two datasets. This is attributed to the principled Bayesian modeling of the underlying FL data generative process in our approaches that can be seen as rigorous generalisation and extension of the existing intuitive algorithms such as FedAvg and FedProx. In particular, the superiority of our methods to the other Bayesian/ensemble approaches verifies the effectiveness of modeling client-wise latent variables θ_i against the

Table 6: (MNIST/FMNIST/EMNIST) Global prediction and personalisation accuracy.

(a) Global prediction performance (initial accuracy)													
Dataset	FL settings			Our Methods		Fed-BABU	Fed-Avg	Fed-Prox	pFedBayes	FedPA	FedBE	FedEM	FedPop
	s	f	τ	NIW	Mix ($K=2$)								
MNIST	5	0.1	1	97.81	97.94	97.16	97.15	97.38	97.32	93.38	97.16	97.38	97.42
			5	95.51	95.68	94.59	94.86	95.28	94.28	94.87	95.11	95.33	94.98
FMNIST	5	0.1	1	84.18	84.28	83.86	81.98	83.25	81.20	81.10	80.35	83.51	82.40
			5	77.48	77.60	76.10	73.70	73.67	72.35	73.47	73.28	76.69	72.93
EMNIST	5	0.2	1	85.40	85.58	84.33	85.27	85.27	84.65	83.20	85.24	85.21	85.27

(b) Personalisation performance													
Dataset	FL settings			Our Methods		Fed-BABU	Fed-Avg	Fed-Prox	pFedBayes	FedPA	FedBE	FedEM	FedPop
	s	f	τ	NIW	Mix ($K=2$)								
MNIST	5	0.1	1	98.78	98.85	97.73	97.83	97.96	97.88	96.66	97.62	97.89	97.91
			5	96.53	96.67	95.89	96.10	96.45	95.49	95.93	96.37	96.46	95.80
FMNIST	5	0.1	1	92.48	92.54	91.03	90.59	90.72	90.31	91.60	89.96	92.10	91.14
			5	89.91	89.53	89.02	86.85	87.35	86.38	88.76	85.00	89.65	86.40
EMNIST	5	0.2	1	88.84	88.97	83.09	87.92	88.39	88.12	85.10	88.37	88.32	88.40

commonly used shared θ modeling. Our methods are especially robust for the scenarios of significant client data heterogeneity, e.g., CIFAR-C-100 personalisation on data with unseen corruption types.

In Table 5, our two models (NIW and Mix.) consistently perform the best (by large margins most of the time) in terms of both global prediction and personalisation for all FL settings. This is attributed to the principled Bayesian modeling of the underlying FL data generative process in our approaches that can be seen as rigorous generalisation and extension of the existing intuitive algorithms such as Fed-Avg and Fed-Prox. In particular, the superiority of our methods to the other Bayesian approach pFedBayes verifies the effectiveness of modeling client-wise latent variables θ_i against the commonly used shared θ modeling, especially for the scenarios of significant client data heterogeneity (e.g., personalisation on data with unseen corruption types).

6.3 Training Dynamics

We visualise the training curves for the global posterior $q(\phi; L_0)$ and for the personalisation training on CIFAR-100 ($s = 100, f = 0.1, \tau = 10$ setting) in Fig. 3. On the right panel (personalisation), we see that the trained model $q(\phi; L_0)$ is improved further by about 4% in test accuracy by our personalisation training. We also visualise the MNIST training curves for our FedHB-NIW and FedHB-Mixture models in Fig. 4 for different client participant rates f . For all different rates, they all converge to the equal (local optimum) loss values, including the centralised optimisation, which supports our theoretical claim on the convergence (Theorem 1). We can also see that increasing the number of participating clients improves the convergence rates, as stated in our Theorem 1, specifically, the effect of N_f on the optimality error formula on the right hand side of (27).

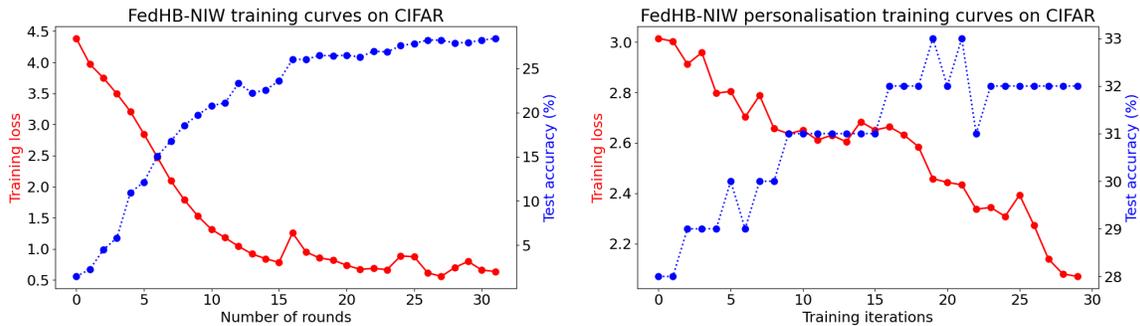


Figure 3: CIFAR-100 training dynamics. (Left) Training curves as FL rounds. (Right) Personalisation training curves. We also superimpose test accuracies.

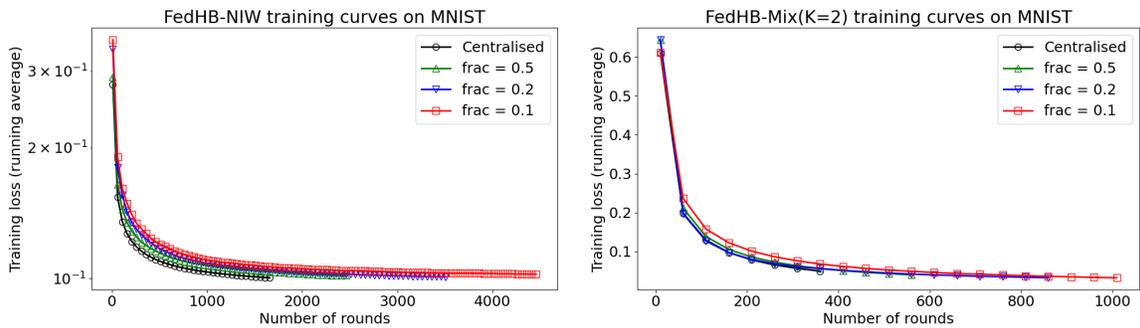


Figure 4: MNIST training convergence with different numbers of participating clients. (Left) NIW and (Right) Mixture ($K = 2$).

6.4 Ablation Study and Hyperparameter Sensitivity

6.4.1 CHANGING MIXTURE ORDER (K) IN OUR MIXTURE MODEL

We test our mixture model with different mixture orders $K = 2, 5, 10$ on CIFAR-100 (Table 7) and CIFAR-C-100 (Table 8). In the last columns of the tables, we also report the performance of the centralised (non-FL) training, in which the batch sampling follows the corresponding FL settings. That is, at each round, the minibatches for SGD (for conventional cross-entropy loss minimisation) are sampled from the data of the participating clients. The centralised training sometimes outperforms the best FL algorithms (our models), but can fail completely especially when data heterogeneity is high (small s) and τ is large. This may be due to overtraining on biased client data for relatively few rounds. Our FL models perform well consistently and stably being comparable to centralised training on its ideal settings (small τ and/or large s).

6.4.2 COMPARISON WITH SIMPLE ENSEMBLE BASELINES.

Our mixture model maintains K backbone networks (specifically $\{r_j\}_{j=1}^K$) where the mixture order K is usually small but greater than 1 (e.g., $K = 2$). Thus it requires extra computational resources than other methods (including our NIW model) that only deal with a single backbone. As a baseline comparison, we aim to come up with some simple extension of

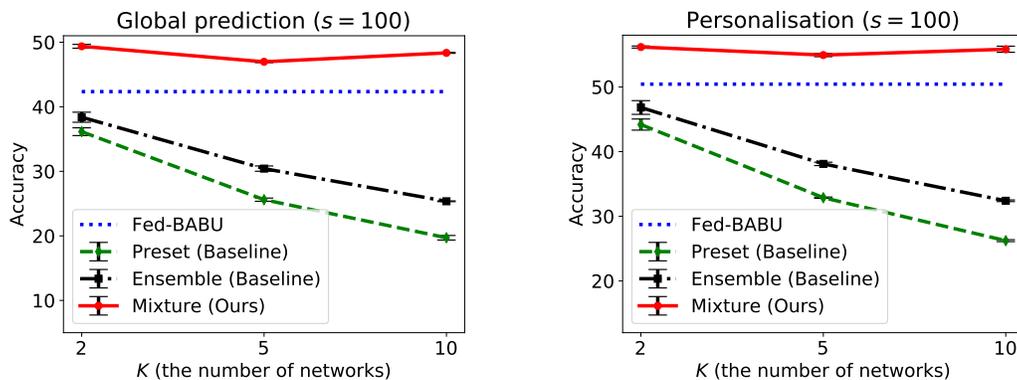
Table 7: (CIFAR-100) Global prediction and personalisation accuracy. Mixture order K varied. Comparison with centralised (non-FL) training.

(a) Global prediction performance (initial accuracy)

FL settings			NIW (Ours)	Mixture (Ours)			FedBABU	Centralised
s	f	τ		$K = 2$	$K = 5$	$K = 10$		
100	0.1	1	49.76 \pm 0.12	49.37 \pm 0.30	46.97 \pm 0.13	48.35 \pm 0.07	42.35 \pm 0.42	52.21 \pm 0.19
		10	29.02 \pm 0.33	29.02 \pm 0.29	29.08 \pm 0.63	29.90 \pm 0.25	27.93 \pm 0.28	36.87 \pm 1.72
	1.0	1	57.80 \pm 0.10	52.94 \pm 0.36	52.24 \pm 0.20	51.92 \pm 0.04	48.17 \pm 0.56	58.50 \pm 0.52
		10	29.53 \pm 0.42	30.55 \pm 0.15	30.85 \pm 0.13	30.24 \pm 0.49	28.67 \pm 0.51	48.16 \pm 0.58
10	0.1	1	37.54 \pm 0.25	38.07 \pm 0.40	39.96 \pm 0.63	39.97 \pm 0.52	35.04 \pm 0.56	28.55 \pm 1.05
		10	18.99 \pm 0.03	18.95 \pm 0.13	18.93 \pm 0.37	18.95 \pm 0.03	18.54 \pm 0.37	9.92 \pm 0.33
	1.0	1	50.40 \pm 0.11	49.52 \pm 0.88	49.82 \pm 0.43	49.51 \pm 0.36	45.41 \pm 0.11	78.37 \pm 0.88
		10	22.87 \pm 0.41	23.59 \pm 0.47	23.69 \pm 0.76	24.28 \pm 1.04	21.92 \pm 0.66	4.40 \pm 0.21

(b) Personalisation performance

FL settings			NIW (Ours)	Mixture (Ours)			FedBABU	Centralised
s	f	τ		$K = 2$	$K = 5$	$K = 10$		
100	0.1	1	54.16 \pm 0.50	56.17 \pm 0.16	54.93 \pm 0.25	55.83 \pm 0.47	50.43 \pm 0.93	53.18 \pm 0.10
		10	36.68 \pm 0.37	36.32 \pm 0.27	37.30 \pm 0.64	37.34 \pm 0.38	35.45 \pm 0.34	38.20 \pm 1.58
	1.0	1	60.36 \pm 0.89	58.82 \pm 0.37	58.16 \pm 0.26	58.32 \pm 0.34	55.87 \pm 0.91	58.49 \pm 0.50
		10	35.92 \pm 0.17	36.22 \pm 0.17	36.44 \pm 0.15	35.91 \pm 0.15	35.58 \pm 0.24	48.17 \pm 0.59
10	0.1	1	79.41 \pm 0.24	79.70 \pm 0.19	79.29 \pm 0.19	77.44 \pm 0.54	75.44 \pm 0.36	42.04 \pm 1.38
		10	67.35 \pm 1.02	67.57 \pm 0.62	67.84 \pm 0.40	67.33 \pm 0.26	66.24 \pm 0.53	17.40 \pm 1.05
	1.0	1	82.71 \pm 0.37	81.03 \pm 0.35	79.91 \pm 0.25	80.92 \pm 0.19	78.92 \pm 0.23	78.43 \pm 0.90
		10	67.78 \pm 1.02	66.74 \pm 0.27	66.50 \pm 0.24	67.30 \pm 0.29	66.25 \pm 0.46	5.13 \pm 0.19


 Figure 5: Comparison between our mixture model and ensemble baselines (K varied) on CIFAR-100.

FedAvg (McMahan et al., 2017) that incorporates multiple (the same K) networks. Here are the detailed descriptions of the baseline ensemble extensions:

Table 8: (CIFAR-C-100) Global prediction and personalisation accuracy. Mixture order K varied. Comparison with centralised (non-FL) training.

(a) Global prediction (initial accuracy) on test splits for the 10 training corruption types								
FL settings			NIW (Ours)	Mixture (Ours)			FedBABU	Centralised
s	f	τ		$K = 2$	$K = 5$	$K = 10$		
100	0.1	1	81.22 \pm 0.14	80.34 \pm 1.44	81.30 \pm 0.21	81.23 \pm 0.79	79.45 \pm 0.71	87.68 \pm 0.48
		4	67.69 \pm 0.74	65.81 \pm 0.84	64.37 \pm 0.99	66.72 \pm 1.46	63.58 \pm 1.28	76.82 \pm 1.17
	1.0	1	91.26 \pm 0.83	86.84 \pm 0.22	86.41 \pm 0.58	87.04 \pm 0.33	86.84 \pm 0.83	90.22 \pm 0.57
		4	73.58 \pm 1.02	74.55 \pm 0.22	74.53 \pm 0.81	75.03 \pm 0.45	71.03 \pm 0.75	87.71 \pm 0.91
50	0.1	1	78.63 \pm 0.39	79.36 \pm 0.24	78.23 \pm 0.33	78.45 \pm 0.40	77.44 \pm 1.17	84.97 \pm 0.22
		4	65.08 \pm 1.75	63.52 \pm 0.48	63.08 \pm 0.55	63.13 \pm 0.21	62.65 \pm 0.12	70.76 \pm 2.93
	1.0	1	89.31 \pm 0.17	88.24 \pm 0.71	87.81 \pm 0.66	87.02 \pm 0.19	86.44 \pm 0.97	89.92 \pm 0.47
		4	70.33 \pm 0.18	70.19 \pm 1.41	71.87 \pm 2.31	74.18 \pm 0.46	67.66 \pm 0.78	86.40 \pm 1.90

(b) Global prediction (initial accuracy) on the original (uncorrupted) CIFAR-100 training sets								
FL settings			NIW (Ours)	Mixture (Ours)			FedBABU	Centralised
s	f	τ		$K = 2$	$K = 5$	$K = 10$		
100	0.1	1	41.55 \pm 0.11	36.99 \pm 0.09	37.75 \pm 0.50	37.78 \pm 0.44	34.76 \pm 0.50	35.10 \pm 0.65
		4	30.84 \pm 0.07	30.60 \pm 0.27	30.51 \pm 0.20	29.13 \pm 0.15	28.31 \pm 0.28	33.64 \pm 1.09
	1.0	1	41.32 \pm 0.32	38.35 \pm 0.73	38.85 \pm 0.96	38.97 \pm 0.26	35.58 \pm 0.41	32.78 \pm 0.22
		4	30.67 \pm 0.12	30.40 \pm 0.44	30.57 \pm 0.28	30.95 \pm 0.28	28.60 \pm 0.28	28.95 \pm 2.09
50	0.1	1	41.04 \pm 0.14	36.41 \pm 0.47	36.03 \pm 0.24	38.32 \pm 0.07	35.44 \pm 0.58	35.71 \pm 0.13
		4	32.29 \pm 0.36	31.50 \pm 0.34	31.19 \pm 0.26	31.06 \pm 0.58	29.68 \pm 0.08	35.28 \pm 1.47
	1.0	1	41.64 \pm 0.21	38.54 \pm 0.42	39.19 \pm 0.37	38.93 \pm 0.15	36.09 \pm 0.28	33.49 \pm 0.17
		4	32.17 \pm 0.48	30.68 \pm 0.46	31.29 \pm 0.65	32.46 \pm 0.20	29.28 \pm 0.17	29.50 \pm 1.24

(c) Personalisation performance on the 9 held-out corruption types								
FL settings			NIW (Ours)	Mixture (Ours)			FedBABU	Centralised
s	f	τ		$K = 2$	$K = 5$	$K = 10$		
100	0.1	1	72.63 \pm 2.13	74.16 \pm 3.04	74.43 \pm 4.94	75.65 \pm 4.68	69.93 \pm 1.24	88.76 \pm 0.31
		4	62.74 \pm 0.94	61.56 \pm 1.69	60.46 \pm 0.13	64.41 \pm 2.66	60.33 \pm 2.12	80.47 \pm 0.99
	1.0	1	83.62 \pm 1.84	84.88 \pm 0.85	83.01 \pm 1.79	82.58 \pm 0.91	77.55 \pm 2.05	87.55 \pm 2.93
		4	64.84 \pm 1.05	67.35 \pm 1.46	64.34 \pm 0.75	65.12 \pm 0.77	53.25 \pm 2.19	76.58 \pm 12.63
50	0.1	1	75.50 \pm 0.74	67.33 \pm 2.83	67.32 \pm 1.35	65.34 \pm 3.19	59.47 \pm 2.48	84.08 \pm 1.03
		4	44.90 \pm 1.23	46.39 \pm 0.83	45.95 \pm 0.37	45.43 \pm 0.34	44.74 \pm 1.30	74.14 \pm 2.51
	1.0	1	81.46 \pm 0.67	81.77 \pm 3.11	78.03 \pm 1.05	74.62 \pm 1.95	67.43 \pm 2.58	87.57 \pm 1.62
		4	50.84 \pm 0.74	48.90 \pm 0.51	45.07 \pm 1.15	48.78 \pm 0.49	40.87 \pm 3.01	81.46 \pm 1.11

1. The server maintains K networks (denoted by $\theta^1, \dots, \theta^K$).
2. We partition the clients into K groups with equal proportions. We will assign each θ^j to each group j ($j = 1, \dots, K$).
3. At each round, each participating client i receives the current model $\theta^{j(i)}$ from the server, where $j(i)$ means the group index to which client i belongs.

Table 9: Running times (in seconds) on CIFAR-100 with ($s = 100, f = 1.0, \tau = 1$) setting.

		Fed-BABU	NIW	Mix. ($K = 2$)	Mix. ($K = 5$)
Train	Client	0.283	0.362	0.481	0.595
	Server	0.027	0.267	1.283	1.434
Global prediction		0.021	0.025	0.040	0.060
Personalisation		1.581	2.158	2.421	2.766

4. The clients perform local updates as usual by warm-start with the received models, and send the updated models back to the server.
5. The server collects updated local models from the clients, and takes the average within each group j to update θ^j .
6. After training, we have trained K networks. At test time, we can use these K networks in two different ways/options: (*Preset* option) Each client i uses the network assigned to its group, i.e., $\theta^{j(i)}$, for both prediction and finetuning/personalisation; (*Ensemble* option) We use all K networks (as an ensemble) for prediction and finetuning.

Note that $K = 1$ exactly reduces to FedAvg (or FedBABU). In Fig. 5 we visualise the performance of these ensemble baselines, compared with our mixture model for different $K = 2, 5, 10$ on CIFAR-100 with ($f = 0.1, \tau = 1$) setting. It clearly shows that these simple ensemble strategies are prone to overfit. The result signifies the importance of the sophisticated negative **log-sum-exp** regularisation in the client/server updates as in (24) and (25) in our mixture model.

6.4.3 HYPERPARAMETER SENSITIVITY

We test sensitivity to some key hyperparameters in our models. For NIW, we have $p_{do} = 1 - p_{drop}$, the MC-dropout probability, where we used $p_{drop} = 0.001$ in the main experiments. In Fig. 2(a) we report the performance of NIW for different values ($p_{drop} = 0, 10^{-4}, 10^{-2}$) on CIFAR-100 with ($s = 100, f = 0.1, \tau = 1$) setting. We see that the performance is not very sensitive to p_{drop} unless it is too large (e.g., 0.01). For the Mixture model, different mixture orders $K = 2, 5, 10$ are contrasted in Fig. 2(b). As seen, having more mixture components does no harm (no overfitting), but we do not see further improvement over $K = 2$ in our experiments (See also results on CIFAR-C-100 in Table 5 in Appendix G).

6.5 Running Times

Although our models achieve significant improvement in prediction accuracy, we have extra computational overhead compared to simpler FL methods like Fed-BABU. To see if this extra cost is allowable, we measure/compare wall clock times in Table 9, where all methods are tested on the same machine, Xeon 2.20GHz CPU with a single RTX 2080 Ti GPU. For NIW, the extra cost in the local client update and personalisation (training) originates from the penalty term in (19), while model weight squaring to compute V_0 in (20) incurs additional cost in server update. For Mixture, the increased time in training is mainly due to the

overhead of computing distances from the K server models in (24) and (25). However, overall the extra costs are not prohibitively large, rendering our methods sufficiently practical.

7. Conclusion

We have proposed a novel hierarchical Bayesian approach to FL where the block-coordinate descent solution to the variational inference leads to a viable algorithm for FL. Our method not only justifies the previous FL algorithms that look intuitive but theoretically less underpinned, but also generalises them even further via principled Bayesian approaches. With strong theoretical support in convergence rate and generalisation error, our approach is also empirically shown to be superior to recent FL approaches by large margin on several benchmarks with various FL settings.

Appendix

Table of Contents

- Detailed derivations (Sec. A)
 - ELBO Derivation for General Framework (Sec. A.1)
 - Normal-Inverse-Wishart (NIW) Model (Sec. A.2)
 - Mixture Model (Sec. A.3)
- Theorems and Proofs (Sec. B)
 - Convergence Analysis (Sec. B.1)
 - Generalisation Error Bound (Sec. B.2)

Appendix A. Detailed Derivations

A.1 ELBO Derivation for General Framework

We derive the ELBO objective (3) for the general Bayesian FL framework. We first derive the KL divergence between the posterior $q(\phi, \theta_{1:N})$ and $p(\phi, \theta_{1:N}|D_{1:N})$ as follows:

$$\text{KL}(q(\phi, \theta_{1:N}) || p(\phi, \theta_{1:N}|D_{1:N})) = \mathbb{E}_q \left[\log \frac{q(\phi) \cdot \prod_i q_i(\theta_i) \cdot p(D_{1:N})}{p(\phi) \cdot \prod_i p(\theta_i|\phi) \cdot \prod_i p(D_i|\theta_i)} \right] \quad (29)$$

$$= \mathbb{E}_{q(\phi)} \left[\log \frac{q(\phi)}{p(\phi)} \right] + \sum_{i=1}^N \mathbb{E}_{q(\phi)q_i(\theta_i)} \left[\log \frac{q_i(\theta_i)}{p(\theta_i|\phi)} \right] + \mathbb{E}_{q(\theta_{1:N})} \left[\log \frac{p(D_{1:N})}{\prod_i p(D_i|\theta_i)} \right] \quad (30)$$

$$= \underbrace{\text{KL}(q(\phi)||p(\phi)) + \sum_{i=1}^N \left(\mathbb{E}_{q(\phi)} [\text{KL}(q_i(\theta_i)||p(\theta_i|\phi))] + \mathbb{E}_{q_i(\theta_i)} [-\log p(D_i|\theta_i)] \right)}_{=: \mathcal{L}(L)} + \log p(D_{1:N}). \quad (31)$$

Since the KL divergence $\text{KL}(q(\phi, \theta_{1:N})||p(\phi, \theta_{1:N}|D_{1:N}))$ is non-negative, so is (31). Hence,

$$\mathcal{L}(L) \geq -\log p(D_{1:N}) \quad (32)$$

That is, $\mathcal{L}(L)$ is an upper bound of the negative data log-likelihood $-\log p(D_{1:N})$, rendering $\mathcal{L}(L)$ our objective function to be minimised. Since $-\mathcal{L}(L)$ is a lower bound of the (log-) evidence, we often call $\mathcal{L}(L)$ the negative evidence lower bound (ELBO) function.

A.2 Normal-Inverse-Wishart (NIW) Model

We define the prior as a conjugate form of Gaussian and Normal-Inverse-Wishart. More specifically, each local client has Gaussian prior $p(\theta_i|\phi) = \mathcal{N}(\theta_i; \mu, \Sigma)$ where $\phi = (\mu, \Sigma)$, and the global latent variable ϕ is distributed as a conjugate prior which is Normal-Inverse-Wishart (NIW),

$$p(\phi) = \mathcal{NIW}(\mu, \Sigma; \Lambda) = \mathcal{N}(\mu; \mu_0, \lambda_0^{-1}\Sigma) \cdot \mathcal{IW}(\Sigma; \Sigma_0, \nu_0), \quad (33)$$

$$p(\theta_i|\phi) = \mathcal{N}(\theta_i; \mu, \Sigma), \quad i = 1, \dots, N, \quad (34)$$

where $\Lambda = \{\mu_0, \Sigma_0, \lambda_0, \nu_0\}$ is the parameters of the NIW. Although Λ can be learned via data marginal likelihood maximisation (e.g., empirical Bayes), but for simplicity we leave it fixed as: $\mu_0 = 0$, $\Sigma_0 = I$, $\lambda_0 = 1$, and $\nu_0 = d + 2$ where d is the number of parameters in θ_i or μ . This choice ensures that the mean of Σ equals I , and μ is distributed as zero-mean Gaussian with covariance Σ .

Next, our choice of the variational density family for $q(\phi)$ is the NIW, not just because it is the most popular parametric family for a pair of mean vector and covariance matrix $\phi = (\mu, \Sigma)$, but it can also admit closed-form expressions in the ELBO function due to the conjugacy as we derive in Sec. A.2.1.

$$q(\phi) := \mathcal{NIW}(\phi; \{m_0, V_0, l_0, n_0\}) = \mathcal{N}(\mu; m_0, l_0^{-1}\Sigma) \cdot \mathcal{IW}(\Sigma; V_0, n_0). \quad (35)$$

Although the scalar parameters l_0, n_0 can be optimised together with m_0, V_0 , their impact is less influential and we find that they make the ELBO optimisation a little bit cumbersome. So we aim to estimate their optimal values in advance with reasonably good quality. To this end, we exploit the conjugacy of the NIW prior-posterior under the Gaussian likelihood. For each θ_i , we pretend that we have instance-wise representative estimates $\theta_i(x, y)$, one for each $(x, y) \in D_i$. For instance, one can view $\theta_i(x, y)$ as the network parameters optimised with the single training instance (x, y) . Then this amounts to observing $|D|$ ($= \sum_{i=1}^N |D_i|$) Gaussian samples $\theta_i(x, y) \sim \mathcal{N}(\theta_i; \mu, \Sigma)$ for $(x, y) \sim D_i$ and $i = 1, \dots, N$. Then applying the NIW conjugacy, the posterior is the NIW with $l_0 = \lambda_0 + |D| = |D| + 1$ and $n_0 = \nu_0 + |D| = |D| + d + 2$. This gives us good approximate estimates for the optimal l_0, n_0 , and we fix them throughout the variational optimisation. Note that this is only heuristics for estimating the scalar parameters l_0, n_0 quickly, and the parameters m_0, V_0 are determined by the principled ELBO optimisation (Sec. A.2.1). That is, $L_0 = \{m_0, V_0\}$. Since the dimension d is large (the number of neural network parameters), we restrict V_0 to be diagonal for computational tractability.

The density family for $q_i(\theta_i)$'s can be a Gaussian, but we find that it is computationally more attractive and numerically more stable to adopt the mixture of two spiky Gaussians that leads to the MC-Dropout (Gal and Ghahramani, 2016). That is,

$$q_i(\theta_i) = \prod_l \left(p_{do} \cdot \mathcal{N}(\theta_i[l]; m_i[l], \epsilon^2 I) + (1 - p_{do}) \cdot \mathcal{N}(\theta_i[l]; 0, \epsilon^2 I) \right), \quad (36)$$

where (i) m_i is the only variational parameters ($L_i = \{m_i\}$), (ii) $\cdot[l]$ indicates the specific column/layer in neural network parameters where l goes over layers and columns of weight matrices, (iii) p_{do} is the (user-specified) hyperparameter where $1 - p_{do}$ corresponds to the

dropout probability, and (iv) ϵ is a tiny constant (e.g., 10^{-6}) that makes two Gaussians spiky, close to the delta function. Now we provide more detailed derivations for the client optimisation and server optimisation.

A.2.1 DETAILED DERIVATIONS FOR NIW MODEL

Client update. We work on the objective function in the general client update optimisation (5). We note that $q(\phi)$ is spiky since our pre-estimated NIW parameters l_0 and n_0 are large (as the entire training data size $|D|$ is added to the initial prior parameters). Due to the spiky $q(\phi)$, we can accurately approximate the second term in (5) as:

$$\mathbb{E}_{q(\phi)}[\text{KL}(q_i(\theta_i)||p(\theta_i|\phi))] \approx \text{KL}(q_i(\theta_i)||p(\theta_i|\phi^*)), \quad (37)$$

where $\phi^* = (\mu^*, \Sigma^*)$ is the mode of $q(\phi)$, which has closed forms for the NIW distribution:

$$\mu^* = m_0, \quad \Sigma^* = \frac{V_0}{n_0 + d + 1}. \quad (38)$$

Note also that a reparametrised sample for (36) can be denoted as \tilde{m}_i , which is defined to be m_i with random $1 - p_{do}$ portion of its parameters set to 0 (i.e., a dropout version of m_i). This allows us to rewrite the first term of (5) as:

$$\mathbb{E}_{q_i(\theta_i)}[-\log p(D_i|\theta_i)] \approx -\log p(D_i|\tilde{m}_i). \quad (39)$$

We can further work out the second term of (5) as follows:

$$\mathbb{E}_{q(\phi)}[\text{KL}(q_i(\theta_i)||p(\theta_i|\phi))] \approx \text{KL}(q_i(\theta_i)||p(\theta_i|\phi^*)) \quad (40)$$

$$= \text{KL}\left(\prod_l \left(p_{do} \cdot \mathcal{N}(\theta_i[l]; m_i[l], \epsilon^2 I) + (1-p_{do}) \cdot \mathcal{N}(\theta_i[l]; 0, \epsilon^2 I)\right) \parallel \mathcal{N}\left(\theta_i; m_0, \frac{V_0}{n_0+d+1}\right)\right) \quad (41)$$

$$= \sum_l \text{KL}\left(p_{do} \cdot \mathcal{N}(\theta_i[l]; m_i[l], \epsilon^2 I) + (1-p_{do}) \cdot \mathcal{N}(\theta_i[l]; 0, \epsilon^2 I) \parallel \mathcal{N}\left(\theta_i[l]; m_0[l], \frac{V_0[l]}{n_0+d+1}\right)\right) \quad (42)$$

$$\approx \sum_l p_{do} \cdot \text{KL}\left(\mathcal{N}(\theta_i[l]; m_i[l], \epsilon^2 I) \parallel \mathcal{N}\left(\theta_i[l]; m_0[l], \frac{V_0[l]}{n_0+d+1}\right)\right) + \sum_l (1-p_{do}) \cdot \text{KL}\left(\mathcal{N}(\theta_i[l]; 0, \epsilon^2 I) \parallel \mathcal{N}\left(\theta_i[l]; m_0[l], \frac{V_0[l]}{n_0+d+1}\right)\right) \quad (43)$$

$$\approx \sum_l p_{do} \cdot \text{KL}\left(\mathcal{N}(\theta_i[l]; m_i[l], \epsilon^2 I) \parallel \mathcal{N}\left(\theta_i[l]; m_0[l], \frac{V_0[l]}{n_0+d+1}\right)\right) + \text{const} \quad (44)$$

$$= p_{do} \cdot \text{KL}\left(\mathcal{N}(\theta_i; m_i, \epsilon^2 I) \parallel \mathcal{N}\left(\theta_i; m_0, \frac{V_0}{n_0+d+1}\right)\right) + \text{const} \quad (45)$$

$$= p_{do} \cdot \frac{1}{2}(n_0 + d + 1)(m_i - m_0)^\top V_0^{-1}(m_i - m_0) + \text{const} \quad (46)$$

In (42) we use the fact that the KL divergence is decomposed into factorised components; In (43) an approximation similar to (Gal and Ghahramani, 2016) is applied, that is,

$\text{KL}(\sum_i \alpha_i \mathcal{N}_i || \mathcal{N}) \approx \sum_i \alpha_i \text{KL}(\mathcal{N}_i || \mathcal{N})$; The second term in (43) is constant with respect to the optimisation variable m_i ; In (45) we apply the decomposable KL divergence property in the reverse direction; Lastly (46) comes from the KL divergence formula for Gaussians. Now combining (39) and (46) allows us to rewrite (5) as:

$$\min_{m_i} \mathcal{L}_i(m_i) := -\log p(D_i | \tilde{m}_i) + \frac{p_{do}}{2} (n_0 + d + 1) (m_i - m_0)^\top V_0^{-1} (m_i - m_0). \quad (47)$$

Also, we use a minibatch version of the first term for a tractable SGD update, which amounts to replacing the first term by the batch average $\mathbb{E}_{(x,y) \sim \text{Batch}}[-\log p(y|x, \tilde{m}_i)]$ while downweighing the second term by the factor of $1/|D_i|$. Note that m_0 and V_0 are fixed during the optimisation. Interestingly (47) generalises the famous FedAvg (McMahan et al., 2017) and FedProx (Li et al., 2018): With $p_{do} = 1$ (i.e., no dropout) and setting $V_0 = \alpha I$ for some constant α , we see that (47) reduces to the client update formula for FedProx where α controls the impact of the proximal term.

Server update. The server optimisation (6) involves two terms, both of which we will show admit closed-form expressions thanks to the conjugacy. Furthermore, we show that the optimal solution (m_0, V_0) of (6) has an analytic form. First, the KL term in (6) is decomposed as:

$$\text{KL}(\mathcal{IW}(\Sigma; V_0, n_0) || \mathcal{IW}(\Sigma; \Sigma_0, \nu_0)) + \mathbb{E}_{\mathcal{IW}(\Sigma; V_0, n_0)}[\text{KL}(\mathcal{N}(\mu; m_0, l_0^{-1}\Sigma) || \mathcal{N}(\mu; \mu_0, \lambda_0^{-1}\Sigma))] \quad (48)$$

By some algebra, (48) becomes identical to the following, up to constant, removing those terms that are not dependent on m_0, V_0 (See Appendix A.2.2 for derivations):

$$\frac{1}{2} \left(n_0 \text{Tr}(\Sigma_0 V_0^{-1}) + \nu_0 \log |V_0| + \lambda_0 n_0 (\mu_0 - m_0)^\top V_0^{-1} (\mu_0 - m_0) \right). \quad (49)$$

Next, the second term of (6) also admits a closed form as follows (Appendix A.2.2 for details):

$$-\mathbb{E}_{q(\phi)q_i(\theta_i)}[\log p(\theta_i | \phi)] = \frac{n_0}{2} \left(p_{do} m_i^\top V_0^{-1} m_i - p_{do} m_0^\top V_0^{-1} m_i - p_{do} m_i^\top V_0^{-1} m_0 + m_0^\top V_0^{-1} m_0 \right. \\ \left. + \frac{1}{n_0} \log |V_0| + \epsilon^2 \text{Tr}(V_0^{-1}) \right) + \text{const}. \quad (50)$$

That is, server's loss function \mathcal{L}_0 , as defined in (6), is the sum of (49) and (50). We can take the gradients of the loss with respect to m_0, V_0 as follows (also plugging $\mu_0 = 0, \Sigma_0 = I, \lambda_0 = 1, \nu_0 = d + 2$):

$$\frac{\partial \mathcal{L}_0}{\partial m_0} = n_0 V_0^{-1} \left((N + 1) m_0 - p_{do} \sum_{i=1}^N m_i \right), \quad (51)$$

$$\frac{\partial \mathcal{L}_0}{\partial V_0^{-1}} = \frac{1}{2} \left(n_0 (1 + N \epsilon^2) I - (N + d + 2) V_0 + n_0 m_0 m_0^\top + n_0 \sum_{i=1}^N \rho(m_0, m_i, p_{do}) \right), \quad (52)$$

$$\text{where } \rho(m_0, m_i, p_{do}) = p_{do} m_i m_i^\top - p_{do} m_0 m_i^\top - p_{do} m_i m_0^\top + m_0 m_0^\top.$$

We set the gradients to zero and solve for them, which yields the optimal solution:

$$m_0^* = \frac{p_{do}}{N+1} \sum_{i=1}^N m_i, \quad V_0^* = \frac{n_0}{N+d+2} \left((1 + N\epsilon^2)I + m_0^*(m_0^*)^\top + \sum_{i=1}^N \rho(m_0^*, m_i, p_{do}) \right). \quad (53)$$

Note that m_i 's are fixed from clients' latest variational parameters.

It is interesting to see that m_0^* in (53) generalises the well-known aggregation step of averaging local models in FedAvg (McMahan et al., 2017) and related methods: when $p_{do} = 1$ (i.e., no dropout), it almost⁶ equals client model averaging. Also, since $\rho(m_0^*, m_i, p_{do} = 1) = (m_i - m_0^*)(m_i - m_0^*)^\top$ when $p_{do} = 1$, we can see that V_0^* in (53) essentially estimates the sample scatter matrix with $(N+1)$ samples, namely clients' m_i 's and server's prior $\mu_0 = 0$, measuring how much they deviate from the center m_0^* . It is known that the dropout can help regularise the model and lead to better generalisation (Gal and Ghahramani, 2016), and with $p_{do} < 1$ our (53) forms a principled optimal solution.

Global prediction. In the inner integral of (8) of the general predictive distribution, we plug $p(\theta|\phi) = \mathcal{N}(\theta; \mu, \Sigma)$ and NIW $q(\phi)$ of (35). This leads to the multivariate Student- t distribution:

$$\int p(\theta|\phi) q(\phi) d\phi = \int \mathcal{N}(\theta; \mu, \Sigma) \cdot \mathcal{NIW}(\phi) d\phi = t_{n_0-d+1} \left(\theta; m_0, \frac{(l_0+1)V_0}{l_0(n_0-d+1)} \right), \quad (54)$$

where $t_\nu(a, B)$ is the multivariate Student- t with location a , scale matrix b , and d.o.f. ν . Then the predictive distribution for a new test input x^* can be estimated as⁷:

$$p(y^*|x^*, D_1, \dots, D_N) = \int p(y^*|x^*, \theta) \cdot t_{n_0-d+1} \left(\theta; m_0, \frac{(l_0+1)V_0}{l_0(n_0-d+1)} \right) d\theta \quad (55)$$

$$\approx \frac{1}{S} \sum_{s=1}^S p(y^*|x^*, \theta^{(s)}), \quad \text{where } \theta^{(s)} \sim t_{n_0-d+1} \left(\theta; m_0, \frac{(l_0+1)V_0}{l_0(n_0-d+1)} \right). \quad (56)$$

Personalisation. With the given personalisation training data D^p , we follow the general framework in (13) to find $v(\theta) \approx p(\theta|D^p, \phi^*)$ in a variational way, where ϕ^* obtained from (38). For the density family for $v(\theta)$ we adopt the same spiky mixture form as (36),

$$v(\theta) = \prod_l \left(p_{do} \cdot \mathcal{N}(\theta[l]; m[l], \epsilon^2 I) + (1 - p_{do}) \cdot \mathcal{N}(\theta[l]; 0, \epsilon^2 I) \right), \quad (57)$$

where m is the variational parameters. This leads to the MC-dropout-like learning objective,

$$\min_m -\log p(D^p|\tilde{m}) + \frac{p_{do}}{2} (n_0 + d + 1) (m - m_0)^\top V_0^{-1} (m - m_0), \quad (58)$$

Once v is trained, our predictive distribution follows the MC sampling (14).

6. Only the constant 1 added to the denominator, which comes from the prior and has the regularising effect.

7. In practice we use a single sample ($S = 1$) for computational efficiency.

A.2.2 MATHEMATICAL DETAILS

The server optimisation (6) in our NIW model involves two terms, both of which we will show admit closed-form expressions thanks to the conjugacy. Furthermore, we show that the optimal solution (m_0, V_0) of (6) has an analytic form. First, the KL term in (6) is decomposed as:

$$\text{KL}(q(\phi)||p(\phi)) = \text{KL}(q(\mu|\Sigma)q(\Sigma) || p(\mu|\Sigma)p(\Sigma)) \quad (59)$$

$$= \mathbb{E}_{q(\Sigma)} \left[\log \frac{q(\Sigma)}{p(\Sigma)} \right] + \mathbb{E}_{q(\Sigma)} \mathbb{E}_{q(\mu|\Sigma)} \left[\log \frac{q(\mu|\Sigma)}{p(\mu|\Sigma)} \right] \quad (60)$$

$$= \underbrace{\text{KL}(\mathcal{IW}(\Sigma; V_0, n_0) || \mathcal{IW}(\Sigma; \Sigma_0, \nu_0))}_{=:kl_a} + \underbrace{\mathbb{E}_{\mathcal{IW}(\Sigma; V_0, n_0)} [\text{KL}(\mathcal{N}(\mu; m_0, l_0^{-1}\Sigma) || \mathcal{N}(\mu; \mu_0, \lambda_0^{-1}\Sigma))]}_{=:kl_b}. \quad (61)$$

First we work on $kl_a = \mathbb{E}_{\mathcal{IW}(\Sigma; V_0, n_0)} [\log \mathcal{IW}(\Sigma; V_0, n_0)] - \mathbb{E}_{\mathcal{IW}(\Sigma; V_0, n_0)} [\log \mathcal{IW}(\Sigma; \Sigma_0, \nu_0)]$. From the definition of Inverse-Wishart (assuming $\Sigma = (d \times d)$),

$$\log \mathcal{IW}(\Sigma; \Psi, \nu) = \frac{\nu}{2} \log |\Psi| - \frac{\nu + d + 1}{2} \log |\Sigma| - \frac{1}{2} \text{Tr}(\Psi \Sigma^{-1}) - \log \Gamma_d(\nu/2) - \frac{\nu d}{2} \log 2, \quad (62)$$

where $\Gamma_d(\cdot)$ is the multivariate Gamma function. We use the following facts from (Bishop, 2006; Braun and McAuliffe, 2008):

$$\mathbb{E}_{\mathcal{IW}(\Sigma; \Psi, \nu)} \log |\Sigma| = -d \log 2 + \log |\Psi| - \sum_{i=1}^d \psi((\nu - i + 1)/2) \quad (63)$$

$$\mathbb{E}_{\mathcal{IW}(\Sigma; \Psi, \nu)} \Sigma^{-1} = \nu \Psi^{-1}, \quad (64)$$

where $\psi(\cdot)$ is the digamma function. Applying these to the terms in kl_a yields:

$$kl_a = \frac{1}{2} \left(n_0 \text{Tr}(\Sigma_0 V_0^{-1}) + \nu_0 \log |V_0| \right) + \text{const (w.r.t. } m_0, V_0). \quad (65)$$

Next, using the closed-form expression for the KL between Gaussians, kl_b becomes:

$$kl_b = \frac{1}{2} \mathbb{E}_{\mathcal{IW}(\Sigma; V_0, n_0)} [\lambda_0 (\mu_0 - m_0)^\top \Sigma^{-1} (\mu_0 - m_0)] + \text{const (w.r.t. } m_0, V_0) \quad (66)$$

$$= \frac{\lambda_0 n_0}{2} (\mu_0 - m_0)^\top V_0^{-1} (\mu_0 - m_0) + \text{const (w.r.t. } m_0, V_0), \quad (67)$$

where in (67) we use the fact (64). Combining (68) and (67), we have:

$$\text{KL}(q(\phi)||p(\phi)) = \frac{1}{2} \left(n_0 \text{Tr}(\Sigma_0 V_0^{-1}) + \nu_0 \log |V_0| + \lambda_0 n_0 (\mu_0 - m_0)^\top V_0^{-1} (\mu_0 - m_0) \right) + \text{const.} \quad (68)$$

Next, we derive the second term of (6) ($=_{utc}$ stands for equality up to constant (w.r.t. m_0, V_0)).

$$\mathbb{E}_{q(\phi)q_i(\theta_i)}[\log p(\theta_i|\phi)] = -\frac{1}{2}\mathbb{E}[\log |\Sigma| + (\theta_i - \mu)^\top \Sigma^{-1}(\theta_i - \mu)] + \text{const (w.r.t. } m_0, V_0) \quad (69)$$

$$=_{utc} -\frac{1}{2}\mathbb{E}_{\mathcal{TW}(\Sigma; V_0, n_0)}[\log |\Sigma|] - \frac{1}{2}\mathbb{E}[(\theta_i - \mu)^\top \Sigma^{-1}(\theta_i - \mu)] \quad (70)$$

$$=_{utc} -\frac{1}{2}\log |V_0| - \frac{1}{2}\text{Tr} \mathbb{E}_{\mathcal{TW}(\Sigma; V_0, n_0)}\left[\Sigma^{-1}\mathbb{E}_{\mathcal{N}(\mu; m_0, l_0^{-1}\Sigma)q_i(\theta_i)}[(\theta_i - \mu)(\theta_i - \mu)^\top]\right] \quad (71)$$

$$=_{utc} -\frac{1}{2}\log |V_0| - \frac{1}{2}\text{Tr} \mathbb{E}_{\mathcal{TW}(\Sigma; V_0, n_0)}\left[\Sigma^{-1}(\rho(m_0, m_i, p_{do}) + \epsilon^2 I + l_0^{-1}\Sigma)\right] \quad (72)$$

$$=_{utc} -\frac{1}{2}\log |V_0| - \frac{1}{2}\text{Tr}\left((\rho(m_0, m_i, p_{do}) + \epsilon^2 I)n_0 V_0^{-1}\right) \quad (73)$$

$$=_{utc} -\frac{n_0}{2}\left(p_{do}m_i^\top V_0^{-1}m_i - p_{do}m_0^\top V_0^{-1}m_i - p_{do}m_i^\top V_0^{-1}m_0 + m_0^\top V_0^{-1}m_0 + \frac{\log |V_0|}{n_0} + \epsilon^2 \text{Tr} V_0^{-1}\right), \quad (74)$$

where $\rho(m_0, m_i, p_{do}) = p_{do}m_i m_i^\top - p_{do}m_0 m_i^\top - p_{do}m_i m_0^\top + m_0 m_0^\top$, and we use the definition of $q_i(\theta_i)$ in (36) and the fact (64).

A.3 Mixture Model

Previously, the NIW model expresses our prior belief where each client i acquires its own network parameters θ_i a priori as a Gaussian-perturbed version of the shared parameters μ , namely $\theta_i|\phi \sim \mathcal{N}(\mu, \Sigma)$, as in (16). This is intuitively appealing, but may not be adequate for capturing more drastic diversity in local data across clients. In the situations where clients' local data distributions, as well as their domains and class label semantics, are highly heterogeneous (possibly even set up for adversarial purpose), it would be ideal to consider *multiple different prototypes* for the network parameters, diverse enough to cover the heterogeneity in data distributions across clients. Motivated from this idea, we introduce a mixture prior model as follows.

First we consider that there are K network parameters (prototypes) that can broadly cover the clients data distributions. They are denoted as high-level latent variables, $\phi = \{\mu_1, \dots, \mu_K\}$, and we let them distributed independently as standard normal a priori,

$$p(\phi) = \prod_{j=1}^K \mathcal{N}(\mu_j; 0, I). \quad (75)$$

We here note some clear distinction from the NIW prior. Whereas the NIW prior (15) only controls the mean μ and covariance Σ in the Gaussian, from which local models θ_i are sampled, the mixture prior (75) is far more flexible in covering highly heterogeneous distributions. Each local model is then assumed to be chosen from one of these K prototypes. Thus the prior distribution for θ_i can be modeled as a mixture,

$$p(\theta_i|\phi) = \sum_{j=1}^K \frac{1}{K} \mathcal{N}(\theta_i; \mu_j; \sigma^2 I), \quad (76)$$

where σ is the hyperparameter that captures perturbation scale, and can be chosen by users or learned from data. Note that we put equal mixing proportions $1/K$ due to the symmetry, a priori. That is, each client can take any of μ_j 's equally likely a priori.

We then describe our choice of the variational density $q(\phi) \prod_i q_i(\theta_i)$ to approximate the posterior $p(\phi, \theta_{1:N} | D_{1:N})$. First, $q_i(\theta_i)$ is chosen as a Gaussian,

$$q_i(\theta_i) = \mathcal{N}(\theta_i; m_i, \epsilon^2 I), \quad (77)$$

with small ϵ . For $q(\phi)$ we consider a Gaussian factorised over μ_j 's, but with small variances, that is,

$$q(\phi) = \prod_{j=1}^K \mathcal{N}(\mu_j; r_j, \epsilon^2 I), \quad (78)$$

where $\{r_j\}_{j=1}^K$ are variational parameters (L_0) and ϵ is small (e.g., 10^{-4}). The main reason why we make $q(\phi)$ spiky is that the resulting near-deterministic $q(\phi)$ allows for computationally efficient and accurate MC sampling during ELBO optimisation as well as test time (global) prediction, avoiding difficult marginalisation (Sec. A.3.1 for details). Although Bayesian inference in general encourages to retain as many plausible latent states as possible under the given evidence (observed data), we aim to model this uncertainty by having many (possibly redundant) prototypes μ_j 's rather than imposing larger variance for a single one (e.g., finite-sample approximation of a smooth distribution).

A.3.1 DETAILED DERIVATIONS FOR MIXTURE MODEL

With the full specification of the prior distribution and the variational density family, we are ready to dig into the client objective function (5) and the server (6).

Client update. Since $q(\phi)$ is spiky, we can accurately approximate the second term of (5) as $\text{KL}(q_i(\theta_i) || p(\theta_i | \phi^*))$ where $\phi^* = \{\mu_j^* = r_j\}_{j=1}^K$ is the mode of $q(\phi)$ since

$$\mathbb{E}_{q(\phi)q_i(\theta_i)}[\log p(\theta_i | \phi)] \approx \mathbb{E}_{q_i(\theta_i)}[\log p(\theta_i | \phi^*)]. \quad (79)$$

Since $q_i(\theta_i)$ is also spiky, $\text{KL}(q_i(\theta_i) || p(\theta_i | \phi^*))$, the KL divergence between a Gaussian and a Gaussian mixture, can be approximated accurately using the single mode sample $m_i \sim q_i(\theta_i)$, that is,

$$\text{KL}(q_i(\theta_i) || p(\theta_i | \phi^*)) \approx \log q_i(m_i) - \log p(m_i | \phi^*) \quad (80)$$

$$= -\log \sum_{j=1}^K \mathcal{N}(m_i; r_j, \sigma^2 I) + \text{const.} = -\log \sum_{j=1}^K \exp\left(-\frac{\|m_i - r_j\|^2}{2\sigma^2}\right) + \text{const.} \quad (81)$$

Note here that we use the fact that m_i disappears in $\log q_i(m_i)$. Plugging it into (5) yields the following optimisation for client i :

$$\min_{m_i} \mathbb{E}_{q_i(\theta_i)}[-\log p(D_i | \theta_i)] - \log \sum_{j=1}^K \exp\left(-\frac{\|m_i - r_j\|^2}{2\sigma^2}\right). \quad (82)$$

It is interesting to see that (82) can be seen as generalisation of FedProx (Li et al., 2018), where the proximal regularisation term in FedProx is extended to *multiple* global models r_j 's, penalizing the local model (m_i) straying away from these prototypes. And if we use a single prototype ($K = 1$), the optimisation (82) exactly reduces to the local update objective of FedProx. Since `log-sum-exp` is approximately equal to `max`, the regularisation term in (82) effectively focuses on the closest global prototype r_j from the current local model m_i , which is intuitively well aligned with our initial modeling motivation, namely each local data distribution is explained by one of the global prototypes. Lastly, we also note that in the SGD optimisation setting where we can only access a minibatch $B \sim D_i$ during the optimisation of (82), we follow the conventional practice: replacing the first term of the negative log-likelihood by a stochastic estimate $\mathbb{E}_{q_i(\theta_i)} \mathbb{E}_{(x,y) \sim B} [-\log p(y|x, \theta_i)]$ and multiplying the second term of regularisation by $\frac{1}{|D_i|}$.

Server update. First, the KL term in (6) can be easily derived as:

$$\text{KL}(q(\phi) \| p(\phi)) = \frac{1}{2} \sum_{j=1}^K \|r_j\|^2 + \text{const.} \quad (83)$$

and the second term of (6) approximated as follows:

$$\mathbb{E}_{q(\phi)q_i(\theta_i)} [\log p(\theta_i | \phi)] \approx \mathbb{E}_{q(\phi)} [\log p(m_i | \phi)] \quad (84)$$

$$\approx \log \sum_{j=1}^K \frac{1}{K} \mathcal{N}(m_i; r_j, \sigma^2 I) \quad (85)$$

$$= \log \sum_{j=1}^K \exp \left(-\frac{\|m_i - r_j\|^2}{2\sigma^2} \right) + \text{const.} \quad (86)$$

where the approximations in (86) become accurate due to spiky $q_i(\theta_i)$ and $q(\phi)$, respectively. Combining the two terms leads to the optimisation problem for the server:

$$\min_{\{r_j\}_{j=1}^K} \frac{1}{2} \sum_{j=1}^K \|r_j\|^2 - \sum_{i=1}^N \log \sum_{j=1}^K \exp \left(-\frac{\|m_i - r_j\|^2}{2\sigma^2} \right). \quad (87)$$

Interestingly, (87) generalises the well-known aggregation step of averaging local models in FedAvg (McMahan et al., 2017) and related methods: Especially when $K = 1$, (87) reduces to quadratic optimisation, admitting the optimal solution $r_1^* = \frac{1}{N+\sigma^2} \sum_{i=1}^N m_i$. The extra term σ^2 in the denominator can be explained by incorporating an extra *zero* local model originating from the prior (interpreted as a *neutral* model) with the discounted weight σ^2 rather than 1.

Although (87) for $K > 1$ can be solved by standard gradient descent, the objective function resembles the (regularised) Gaussian mixture log-likelihood, and we can apply the Expectation-Maximisation (EM) algorithm (Dempster et al., 1977) instead. Using Jensen's bound with convexity of the negative `log` function, we have the following alternating steps⁸:

8. Although one can perform several EM steps until convergence, in practice, we find that only one EM step per round is sufficient.

- **E-step:** With the current $\{r_j\}_{j=1}^K$ fixed, compute the prototype assignment probabilities for each local model m_i :

$$c(j|i) = \frac{k_{ij}}{\sum_{j=1}^K k_{ij}}, \quad \text{where } k_{ij} = \exp\left(-\frac{\|m_i - r_j\|^2}{2\sigma^2}\right). \quad (88)$$

- **M-step:** With the current assignments $c(j|i)$ fixed, we solve:

$$\min_{\{r_j\}} \frac{1}{2} \sum_j \|r_j\|^2 + \frac{1}{2\sigma^2} \sum_{i,j} c(j|i) \cdot \|m_i - r_j\|^2, \quad (89)$$

which admits the closed form solution:

$$r_j^* = \frac{\frac{1}{N} \sum_{i=1}^N c(j|i) \cdot m_i}{\frac{\sigma^2}{N} + \frac{1}{N} \sum_{i=1}^N c(j|i)}, \quad j = 1, \dots, K. \quad (90)$$

The server update equation (90) has intuitive meaning that the new prototype r_j becomes the *weighted* average of the local models m_i 's where the weights $c(j|i)$ are determined by the proximity to r_j (i.e., those m_i 's that are closer to r_j have more contribution, and vice versa). This can be seen as an extension of the aggregation step in FedAvg to the multiple prototype case.

Global prediction. By plugging the mixture prior $p(\theta|\phi)$ of (76) and the factorised spiky Gaussian $q(\phi)$ of (78) into the inner integral of (8), we have predictive distribution averaged equally over $\{r_j\}_{j=1}^K$ approximately, that is, $\int p(\theta|\phi) q(\phi) d\phi \approx \frac{1}{K} \sum_{j=1}^K p(y^*|x^*, r_j)$. Unfortunately this is not ideal for our original intention where only one specific model r_j out of K candidates is dominantly responsible for the local data. To meet this intention, we extend our model so that the input point x^* can affect θ together with ϕ , and with this modification our predictive probability can be derived as:

$$p(y^*|x^*, D_{1:N}) = \iint p(y^*|x^*, \theta) p(\theta|x^*, \phi) p(\phi|D_{1:N}) d\theta d\phi \quad (91)$$

$$\approx \iint p(y^*|x^*, \theta) p(\theta|x^*, \phi) q(\phi) d\theta d\phi \quad (92)$$

$$\approx \int p(y^*|x^*, \theta) p(\theta|x^*, \{r_j\}_{j=1}^K) d\theta. \quad (93)$$

To deal with the tricky part of inferring $p(\theta|x^*, \{r_j\}_{j=1}^K)$, we introduce a fairly practical strategy of fitting a *gating function*. The idea is to regard $p(\theta|x^*, \{r_j\}_{j=1}^K)$ as a mixture of experts (Jacobs et al., 1991; Jordan and Jacobs, 1994) where the prototypes r_j 's serving as experts,

$$p(\theta|x^*, \{r_j\}_{j=1}^K) := \sum_{j=1}^K g_j(x^*) \cdot \delta(\theta - r_j), \quad (94)$$

where $\delta(\cdot)$ is the Dirac's delta function, and $g(x)$ is a gating function that outputs a K -dimensional softmax vector. Intuitively, the gating function determines which of the K

prototypes $\{r_j\}_{j=1}^K$ the model θ for the test point x^* belongs to. With (94), the predictive probability in (93) is written as:

$$p(y^*|x^*, D_{1:N}) \approx \sum_{j=1}^K g_j(x^*) \cdot p(y^*|x^*, r_j). \quad (95)$$

However, since we do not have this oracle $g(x)$, we introduce and fit a neural network to the local training data during the training stage. Let $g(x; \beta)$ be the gating network with the parameters β . To train it, we follow the FedAvg⁹ strategy. In the client update stage at each round, while we update the local model m_i with a minibatch $B \sim D_i$, we also find the prototype closest to m_i , namely $j^* := \arg \min_j \|m_i - r_j\|$. Then we form another minibatch of samples $\{(x, j^*)\}_{x \sim B}$ (input x and class label j^*), and update $g(x; \beta)$ by SGD. The updated (local) β 's from the clients are then aggregated (by simple averaging) by the server, and distributed back to the clients as an initial iterate for the next round.

Personalisation. For $p(\theta|D^p, \phi^*)$ in the general framework (12), we define the variational distribution $v(\theta) \approx p(\theta|D^p, \phi^*)$ as:

$$v(\theta) = \mathcal{N}(\theta; m, \epsilon^2 I), \quad (96)$$

where ϵ is small positive constant, and m is the only parameters that we learn. Our personalisation training amounts to ELBO optimisation for $v(\theta)$ as in (13), which reduces to:

$$\min_m \mathbb{E}_{v(\theta)}[-\log p(D^p|\theta)] - \log \sum_{j=1}^K \exp\left(-\frac{\|m - r_j\|^2}{2\sigma^2}\right). \quad (97)$$

Once we have optimal m (i.e., $v(\theta)$), our predictive model becomes:

$$p(y^p|x^p, D^p, D_{1:N}) \approx p(y^p|x^p, m), \quad (98)$$

which is done by feed-forwarding test input x^p through the network deployed with the parameters m .

Appendix B. Theorems and Proofs

B.1 Convergence Analysis

Our (general) FL algorithm is a special block-coordinate SGD optimisation of the ELBO function (3) with respect to the $(N+1)$ parameter groups: L_0 (of $q(\phi; L_0)$), L_1 (of $q_1(\theta_1; L_1)$), \dots , and L_N (of $q_N(\theta_N; L_N)$). In this section we will provide a theorem that guarantees convergence of the algorithm to a local minimum of the ELBO objective function under some mild assumptions. We will also analyse the convergence rate. Note that although our FL algorithm is a special case of the general block-coordinate SGD optimisation, we may not directly apply the existing convergence results for the regular block-coordinate SGD methods

9. We also follow the FedBABU (Oh et al., 2022) strategy by updating only the body of β and fixing/sharing the random classification head across the server and clients.

since they mostly rely on non-overlapping blocks with cyclic or uniform random block selection strategies (Beck and Tetruashvili, 2013; Wang and Banerjee, 2014). As the block selection strategy in our FL algorithm is unique with overlapping blocks and non-uniform random block selection, we provide our own analysis here. Promisingly, we show that in accordance with general regular block-coordinate SGD (cyclic/uniform non-overlapping block selection), our FL algorithm has $O(1/\sqrt{t})$ convergence rate, which is also asymptotically the same as that of the conventional (holistic, non-block-coordinate) SGD optimisation. Note that this section is about the convergence of our algorithm to an (local) optimum of the training objective (ELBO). The question of how well this optimal model trained on empirical data performs on the unseen data points will be discussed in Sec. B.2.

First we formally describe our FL algorithm as a block-coordinate SGD optimisation. For ease of exhibition, we will simplify the notation: The objective function in (3) is denoted as $f(x)$ where $x = [x_0, x_1, \dots, x_N]$ is the optimisation variables corresponding to $x_0 := L_0$, $x_1 := L_1$, \dots , $x_N := L_N$. That is, x_0 is server's parameters while x_i ($i = 1, \dots, N$) is worker i 's parameters. We let x_u be the partial vector of x selected by the index set $u \subseteq \{0, 1, \dots, N\}$, and x_{-u} be the vector of the rest elements. Similarly $\nabla_u f(x)$ indicates the gradient vector with only elements at the indices in u . Let x^t be the iterate at iteration t . Our FL algorithm is formally defined in Alg. 10.

Now we state our convergence theorem. We first need the following mild assumptions:

Assumption 1. Our objective function $f(x)$ is *locally* convex, and $f_t(x)$ is also *locally* convex for all iterations t , where f_t is the minibatch version of f defined on the minibatch data $batch_t$ (so that $\mathbb{E}_{batch_t}[f_t(x)] = f(x)$). Actually, the latter implies the former. Although the negative ELBO is in general non-convex globally, we can regard it as a convex function within a local neighborhood, as is usually assumed in non-convex analysis (Bertsekas, 2016) and other FL analysis (Li et al., 2020).

Assumption 2. For all t , $f_t(x)$ has Lipschitz continuous gradient with constant \bar{L} .

Assumption 3. For all t , $\|\nabla f_t(x)\| \leq R_f$ and $\|x - x'\| \leq D$ for any x, x' , where R_f and D are some constants.

Theorem 1 (Convergence analysis). *Let $\eta_t = \bar{L} + \sqrt{t}$, $M_S = \frac{M \cdot N_f}{N}$, and $\bar{x}^T = \frac{1}{T} \sum_{t=1}^T x^t$ by following our FL algorithm. With Assumptions 1-3, the following holds for any T :*

$$\mathbb{E}[f(\bar{x}^T)] - f(x^*) \leq \frac{N + N_f}{N_f} \cdot \frac{\frac{\sqrt{T} + \bar{L}}{2} D^2 + R_f^2 \sqrt{T}}{T} = O\left(\frac{1}{\sqrt{T}}\right), \quad (101)$$

where x^* is the (local) optimum, and the expectation is taken over randomness in minibatches and block selections $\{u_t\}_{t=0}^{T-1}$.

Remark. *Theorem 1 states that \bar{x}^t converges to the optimal point x^* in expectation at the rate of $O(1/\sqrt{t})$. This convergence rate asymptotically equals that of the conventional (holistic, non-block-coordinate) SGD algorithm.*

To prove the theorem, we note that the algorithm Alg. 10 overall repeats the following three steps per round: i) sample the subset of clients u_t from $\{1, \dots, N\}$ with $|u_t| = N_f$, ii) update x_{u_t} for M iterations, and iii) update x_0 for M_S iterations. Thus in the long-term

Algorithm 10 Bayesian FL Algorithm as Block-Coordinate Descent.

We define the following hyperparameters:

- N_f ($\leq N$) = the number of participating clients at each round.
- M = the number of SGD iterations per round for updating the (participating) clients.
- Let M_S = the number of SGD iterations per round for updating the server.
- Let η_t = the reciprocal of the SGD learning rate at iteration t .

Initialise the global iteration counter $t = 0$.

for each round **do**

- Select N_f clients uniformly at random from $\{1, \dots, N\}$ without replacement. Let $u_t \subseteq \{1, \dots, N\}$ be the set of the participants ($|u_t| = N_f$).

- (Client update) For each of M iterations,

1. Perform an SGD update for the block u_t . That is,

$$x^{t+1} := [x_{u_t}^{t+1}; x_{-u_t}^t], \quad \text{where } x_{u_t}^{t+1} = x_{u_t}^t - \frac{1}{\eta_t} \nabla_{u_t} f_t(x^t), \quad (99)$$

where f_t is the minibatch version of f defined on the minibatch data $batch_t$ (so that $\mathbb{E}_{batch_t}[f_t(x)] = f(x)$). Note that this update is actually done independently over the participating clients $i \in u_t$ due to the separable objective, i.e., from (4) to (5).

2. $t \leftarrow t + 1$.

- (Server update) For each of M_S iterations,

1. Perform SGD update for the index (singleton block) 0. That is,

$$x^{t+1} := [x_0^{t+1}; x_{-0}^t], \quad \text{where } x_0^{t+1} = x_0^t - \frac{1}{\eta_t} \nabla_0 f_t(x^t). \quad (100)$$

2. $t \leftarrow t + 1$.

end for

view, we can see that the algorithm proceeds as follows: At each iteration t , we select u_t as

$$u_t = \begin{cases} \{0\} & \text{with prob. } \frac{M_S}{M+M_S} \\ \text{Size-}N_f \text{ subset uniformly from } \{1, \dots, N\} & \text{with prob. } \frac{M}{M+M_S} \end{cases}, \quad (102)$$

and update the iterate as

$$x^{t+1} := [x_{u_t}^{t+1}; x_{-u_t}^t], \quad \text{where } x_{u_t}^{t+1} = x_{u_t}^t - \frac{1}{\eta_t} \nabla_{u_t} f_t(x^t). \quad (103)$$

We will use this long-term view strategy in our proof. Next we state the following lemma that is motivated from (Wang and Banerjee, 2014), which is useful in our proof.

Lemma 3. *Assume $\eta_t > \bar{L}$, and f_t is (locally) convex with Lipschitz continuous gradient with constant \bar{L} . For any subset $u \subseteq \{0, 1, \dots, N\}$, we define x_{t+1} as:*

$$x_u^{t+1} = x_u^t - \frac{1}{\eta_t} \nabla_u f_t(x^t) \quad \text{and} \quad x_{-u}^{t+1} = x_{-u}^t. \quad (104)$$

Then the following holds for any x :

$$\langle \nabla_u f_t(x^t), x_u^t - x_u \rangle \leq \frac{\eta_t}{2} (\|x - x^t\|^2 - \|x - x^{t+1}\|^2) + \frac{R_f^2}{2(\eta_t - \bar{L})}. \quad (105)$$

Proof [Proof of Lemma 3] By definition, $\nabla_u f_t(x^t) + \eta_t(x_u^{t+1} - x_u^t) = 0$. Then for any x , we have

$$\langle \nabla_u f_t(x^t), x_u^{t+1} - x_u \rangle = -\eta_t \langle x_u^{t+1} - x_u^t, x_u^{t+1} - x_u \rangle \quad (106)$$

$$= \frac{\eta_t}{2} (\|x_u - x_u^t\|^2 - \|x_u - x_u^{t+1}\|^2 - \|x_u^{t+1} - x_u^t\|^2) \quad (107)$$

$$= \frac{\eta_t}{2} (\|x - x^t\|^2 - \|x - x^{t+1}\|^2 - \|x_u^{t+1} - x_u^t\|^2). \quad (108)$$

Note that (108) follows from (107) since x^t and x^{t+1} only differ at indices u . Since f_t has Lipschitz continuous gradient,

$$f_t(x^{t+1}) - f_t(x^t) \leq \langle \nabla_u f_t(x^t), x_u^{t+1} - x_u^t \rangle + \frac{\bar{L}}{2} \|x_u^{t+1} - x_u^t\|^2 \quad (109)$$

$$= \langle \nabla_u f_t(x^t), x_u^{t+1} - x_u \rangle + \frac{\bar{L}}{2} \|x_u^{t+1} - x_u^t\|^2 - \langle \nabla_u f_t(x^t), x_u^t - x_u \rangle \quad (110)$$

$$= \frac{\eta_t}{2} (\|x - x^t\|^2 - \|x - x^{t+1}\|^2 - \|x_u^{t+1} - x_u^t\|^2) + \frac{\bar{L}}{2} \|x_u^{t+1} - x_u^t\|^2 - \langle \nabla_u f_t(x^t), x_u^t - x_u \rangle, \quad (111)$$

where we plugged (108) into (110). We rearrange the terms in (111) as follows:

$$\langle \nabla_u f_t(x^t), x_u^t - x_u \rangle \leq \frac{\eta_t}{2} (\|x - x^t\|^2 - \|x - x^{t+1}\|^2) + \frac{\bar{L} - \eta_t}{2} \|x_u^{t+1} - x_u^t\|^2 + (f_t(x^t) - f_t(x^{t+1})). \quad (112)$$

Due to the convexity of f_t , we can bound the last term in (112) as

$$f_t(x^t) - f_t(x^{t+1}) \leq \langle \nabla f_t(x^t), x^t - x^{t+1} \rangle \quad (113)$$

$$= \langle \nabla_u f_t(x^t), x_u^t - x_u^{t+1} \rangle \quad (114)$$

$$\leq \frac{1}{2\alpha} \|\nabla_u f_t(x^t)\|^2 + \frac{\alpha}{2} \|x_u^t - x_u^{t+1}\|^2 \quad (\text{for any } \alpha > 0). \quad (115)$$

Plugging (115) into (112) and choosing $\alpha = \eta_t - \bar{L} (> 0)$ yields:

$$\langle \nabla_u f_t(x^t), x_u^t - x_u \rangle \leq \frac{\eta_t}{2} (\|x - x^t\|^2 - \|x - x^{t+1}\|^2) + \frac{\|\nabla_u f_t(x^t)\|^2}{2(\eta_t - \bar{L})}. \quad (116)$$

Applying Assumption 3 of the bounded gradient norm completes the proof. \blacksquare

Now we are ready to prove our convergence theorem (Theorem 1).

Proof [Proof of Theorem 1] We first aim to bound $\langle \nabla f(x^t), x^t - x \rangle$, as it upper-bounds of $f(x^t) - f(x)$ for convex f . Note that by conditioning on x^t , we can only deal with randomness in minibatch at t ($batch_t$), that is,

$$\langle \nabla f(x^t), x^t - x \rangle = \mathbb{E}_{batch_t}[\langle \nabla f_t(x^t), x^t - x \rangle]. \quad (117)$$

Further conditioning on $batch_t$ leads to:

$$\langle \nabla f_t(x^t), x^t - x \rangle = \langle \nabla_0 f_t(x^t), x_0^t - x_0 \rangle + \sum_{i=1}^N \langle \nabla_i f_t(x^t), x_i^t - x_i \rangle. \quad (118)$$

Here we aim to rewrite the summation term in (118) in terms of size N_f blocks. To this end, let us consider:

$$\sum_{u \in 2^{N, N_f}} \langle \nabla_u f_t(x^t), x_u^t - x_u \rangle, \quad (119)$$

where $2^{N, N_f}$ is defined as *the set of all subsets of $\{1, \dots, N\}$ with size N_f and no repeating elements*. For instance, $2^{5, 3}$ contains $\{1, 3, 4\}$ and $\{2, 4, 5\}$, among others. Obviously $|2^{N, N_f}| = \binom{N}{N_f}$, and each particular index $i \in \{1, \dots, N\}$ appears exactly $\binom{N-1}{N_f-1}$ times in the sum (119). Thus we can establish the following identity:

$$\sum_{u \in 2^{N, N_f}} \langle \nabla_u f_t(x^t), x_u^t - x_u \rangle = \binom{N-1}{N_f-1} \sum_{i=1}^N \langle \nabla_i f_t(x^t), x_i^t - x_i \rangle. \quad (120)$$

We plug (120) into (118) and apply Lemma 3:

$$\begin{aligned} \langle \nabla f_t(x^t), x^t - x \rangle &= \langle \nabla_0 f_t(x^t), x_0^t - x_0 \rangle + \frac{1}{\binom{N-1}{N_f-1}} \sum_{u \in 2^{N, N_f}} \langle \nabla_u f_t(x^t), x_u^t - x_u \rangle \\ &\leq \frac{\eta_t}{2} (\|x - x^t\|^2 - \|x - x^{t+1}(0, batch_t)\|^2) + \frac{R_f^2}{2(\eta_t - \bar{L})} + \\ &\quad \frac{1}{\binom{N-1}{N_f-1}} \sum_{u \in 2^{N, N_f}} \left(\frac{\eta_t}{2} (\|x - x^t\|^2 - \|x - x^{t+1}(u, batch_t)\|^2) + \frac{R_f^2}{2(\eta_t - \bar{L})} \right), \end{aligned} \quad (121)$$

where we define $x^{t+1}(u, batch_t)$ for any subset $u \subseteq \{0, 1, \dots, N\}$ as: $x_u^{t+1} := x_u^t - (1/\eta_t) \nabla_u f_t(x^t)$ and $x_{-u}^{t+1} := x_{-u}^t$. Although we can simply use x^{t+1} , here we use this explicit notation to emphasise dependency of x^{t+1} on i and $batch_t$. By letting $g(x, x^t, x^{t+1}) := \frac{\eta_t}{2} (\|x - x^t\|^2 - \|x - x^{t+1}\|^2) + \frac{R_f^2}{2(\eta_t - \bar{L})}$, we can express (122) succinctly to yield:

$$\langle \nabla f_t(x^t), x^t - x \rangle \leq g(x, x^t, x^{t+1}(0, batch_t)) + \frac{1}{\binom{N-1}{N_f-1}} \sum_{u \in 2^{N, N_f}} g(x, x^t, x^{t+1}(u, batch_t)). \quad (123)$$

For the second term, we use the uniform expectation to replace the sum (i.e., $\sum_{u \in 2^{N, N_f}} \psi(u) = \binom{N}{N_f} \mathbb{E}_{u \sim 2^{N, N_f}}[\psi(u)]$ for any function ψ). Using $\binom{N}{N_f} / \binom{N-1}{N_f-1} = N/N_f$,

$$\langle \nabla f_t(x^t), x^t - x \rangle \leq g(x, x^t, x^{t+1}(0, \text{batch}_t)) + \frac{N}{N_f} \mathbb{E}_{u \sim 2^{N, N_f}}[g(x, x^t, x^{t+1}(u, \text{batch}_t))], \quad (124)$$

and the right hand side of (124) can be written as:

$$\frac{M + M_S}{M_S} \left(\frac{M_S}{M + M_S} g(x, x^t, x^{t+1}(0, \text{batch}_t)) + \frac{M}{M + M_S} \mathbb{E}_{u \sim 2^{N, N_f}}[g(x, x^t, x^{t+1}(u, \text{batch}_t))] \right), \quad (125)$$

where we use our specification of $M_S = \frac{M \cdot N_f}{N}$. Note that the expression inside the parentheses is exactly the expectation of $g(x, x^t, x^{t+1}(u_t, \text{batch}_t))$ over the random index set u_t following our client-server selection strategy in the long term view, that is, (102). Then we have the following result:

$$\langle \nabla f_t(x^t), x^t - x \rangle \leq \frac{M + M_S}{M_S} \mathbb{E}_{u_t}[g(x, x^t, x^{t+1}(u_t, \text{batch}_t))], \quad (126)$$

where u_t follows (102).

As we have conditioned all quantities on batch_t , we now take the expectation over batch_t .

$$\langle \nabla f(x^t), x^t - x \rangle = \mathbb{E}_{\text{batch}_t}[\langle \nabla f_t(x^t), x^t - x \rangle] \quad (127)$$

$$\leq \frac{M + M_S}{M_S} \mathbb{E}_{\text{batch}_t, u_t} \left[\frac{\eta_t}{2} (\|x - x^t\|^2 - \|x - x^{t+1}\|^2) + \frac{R_f^2}{2\sqrt{t}} \right], \quad (128)$$

where we drop the dependency in $x^{t+1}(u_t, \text{batch}_t)$ in notation, and use $\eta_t = \bar{L} + \sqrt{t}$. Since f is convex,

$$f(x^t) - f(x) \leq \langle \nabla f(x^t), x^t - x \rangle, \quad (129)$$

and taking the expectation over x^t leads to:

$$\mathbb{E}[f(x^t)] - f(x) \leq \frac{M + M_S}{M_S} \left(\frac{\eta_t}{2} (\mathbb{E}\|x - x^t\|^2 - \mathbb{E}\|x - x^{t+1}\|^2) + \frac{R_f^2}{2\sqrt{t}} \right). \quad (130)$$

By telescoping ($\frac{1}{T} \sum_{t=1}^T$) and using $\frac{M + M_S}{M_S} = \frac{N + N_f}{N_f}$, we have:

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T f(x^t) \right] - f(x) \leq \frac{N + N_f}{N_f} \frac{1}{T} \left(\frac{1}{2} \sum_{t=1}^T \eta_t (\mathbb{E}\|x - x^t\|^2 - \mathbb{E}\|x - x^{t+1}\|^2) + R_f^2 \sum_{t=1}^T \frac{1}{2\sqrt{t}} \right). \quad (131)$$

There are two sums in the right hand side of (131), and they can be bounded succinctly as follows. First, we use the simple calculus to bound the second sum: $\sum_{t=1}^T \frac{1}{2\sqrt{t}} \leq \int_1^T \frac{1}{2\sqrt{z}} dz + \frac{1}{2} \leq \sqrt{T}$. Next, we let $a_t := \mathbb{E}\|x - x^t\|^2$, and the first sum is written as:

$\eta_1(a_1 - a_2) + \dots + \eta_T(a_T - a_{T+1}) = \sum_{t=1}^T (\eta_t - \eta_{t-1})a_t - \eta_T a_{T+1}$ by letting $\eta_0 = 0$. Using $a_t \leq D^2$ from Assumption 3, this sum is bounded above by $D^2 \sum_{t=1}^T (\eta_t - \eta_{t-1}) = D^2(\eta_T - \eta_0) = (\sqrt{T} + \bar{L})D^2$. Plugging these bounds into (131) and applying Jensen's inequality to the left hand side (i.e., $\mathbb{E}[(1/T) \sum_{t=1}^T f(x^t)] \geq \mathbb{E}[f((1/T) \sum_{t=1}^T x^t)] = \mathbb{E}[f(\bar{x}^T)]$) yields:

$$\mathbb{E}[f(\bar{x}^T)] - f(x) \leq \frac{N + N_f}{N_f} \cdot \frac{\frac{\sqrt{T} + \bar{L}}{2} D^2 + R_f^2 \sqrt{T}}{T}, \quad (132)$$

for any x , which completes the proof. \blacksquare

B.2 Generalisation Error Bound

In this section we will discuss generalisation performance of our proposed algorithm, answering the question of how well the Bayesian FL model trained on empirical data performs on the unseen data points. We aim to provide the upper bound of the generalisation error averaged over the posterior distribution of the model parameters $(\phi, \{\theta_i\}_{i=1}^N)$, by linking it to the expected empirical error with some additional complexity terms.

We begin with the regression-based data modeling perspective and related assumption-s/notations. We denote by $P^i(x, y)$ the *true* data distribution for client i ($i = 1, \dots, N$). We assume that the target y is real vector-valued ($y \in \mathbb{R}^{S_y}$), and there exists a *true regression function* $f^i : \mathbb{R}^{S_x} \rightarrow \mathbb{R}^{S_y}$ for each i . That is,

$$P^i(y|x) = \mathcal{N}(y; f^i(x), \sigma_\epsilon^2 I), \quad (133)$$

where σ_ϵ^2 is constant Gaussian output noise variance. Let $D_i = (X^i, Y^i) \sim P^i(x, y)$ be the i.i.d. training data of size n for each client i .

In our FL model, we assume that our backbone network is an MLP with L hidden layers of width M , and all activation functions $\sigma(\cdot)$ are Lipschitz continuous with constant 1. The parameters θ of the MLP are also assumed to be bounded, more formally, the parameter space Θ is defined as:

$$\theta \in \Theta = \{\theta \in \mathbb{R}^G : \|\theta\|_\infty \leq B, \text{MLP with } L \text{ layers of width } M\}. \quad (134)$$

Note that $G = \dim(\theta)$ and B is the maximal norm bound. The MLP defines a regression function $f_\theta : \mathbb{R}^{S_x} \rightarrow \mathbb{R}^{S_y}$, and the θ -induced predictive distribution is denoted as:

$$P_\theta(y|x) = \mathcal{N}(y; f_\theta(x), \sigma_\epsilon^2 I), \quad (135)$$

where we assume that the true noise variance is known.

For notational convenience, we denote by $f(X^i)$ the concatenated vector of $f(x)$ for all $x \in X^i$, i.e., $f(X^i) = [f(x)]_{x \in X^i}$, where $f(\cdot)$ is either the true $f^i(\cdot)$ or the model $f_\theta(\cdot)$. Extending this notation, simply writing f^i or f_θ means infinite-dimensional (population) responses, that is, $f^i = [f^i(x)]_{x \in \mathbb{R}^{S_x}}$ and f_θ similarly. For instance, $\|f_\theta - f^i\|_\infty$ stands for the worst-case difference, namely $\max_{x \in \mathbb{R}^{S_x}} \|f_\theta(x) - f^i(x)\|$. As a generalisation error measure,

we consider the *expected squared Hellinger distance* between the true P^i and the model P_θ . Formally,

$$d^2(P_\theta, P^i) = \mathbb{E}_{x \sim P^i(x)} [H^2(P_\theta(y|x), P^i(y|x))] = \mathbb{E}_{x \sim P^i(x)} \left[1 - \exp \left(- \frac{\|f_\theta(x) - f^i(x)\|_2^2}{8\sigma_\epsilon^2} \right) \right]. \quad (136)$$

More specifically, we will bound the posterior-averaged distance $\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{q_i^*(\theta_i)} [d^2(P_{\theta_i}, P^i)]$, where $\{q_i^*(\theta_i)\}_{i=1}^N$ is an optimal solution of our FL-ELBO optimisation problem¹⁰ (We showed in Sec. B.1 that our block-coordinate FL algorithm converges to this optimal solution in $O(1/\sqrt{t})$ rate).

Theorem 2 (Generalisation error analysis). *Assume that the variational density family for $q_i(\theta_i)$ is rich enough to subsume Gaussian. The optimal solution $(\{q_i^*(\theta_i)\}_{i=1}^N, q^*(\phi))$ of our FL-ELBO optimisation problem (3) satisfies (with high probability):*

$$\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{q_i^*(\theta_i)} [d^2(P_{\theta_i}, P^i)] \leq O\left(\frac{1}{n}\right) + C \cdot \epsilon_n^2 + C' \left(r_n + \frac{1}{N} \sum_{i=1}^N \lambda_i^* \right), \quad (137)$$

where $C, C' > 0$ are constant, $\lambda_i^* = \min_{\theta \in \Theta} \|f_\theta - f^i\|_\infty^2$ is the best error within our backbone Θ , r_n is defined as

$$r_n = \frac{G(L+1)}{n} \log M + \frac{G}{n} \log \left(S_x \sqrt{\frac{n}{G}} \right), \quad (138)$$

S_x is the dimensionality of input x , G is the dimensionality of θ , and $\epsilon_n = \sqrt{r_n} \log^\delta(n)$ for $\delta > 1$ constant.

Remark. *Theorem 2 implies that the optimal solution for our FL-ELBO optimisation problem (attainable by our block-coordinate FL algorithm) is asymptotically optimal, since the right hand side of (137) converges to 0 as the training data size $n \rightarrow \infty$. This is easy to verify: as $n \rightarrow \infty$, $r_n \rightarrow 0$ obviously, accordingly $\epsilon_n \rightarrow 0$, and the last term $\frac{1}{N} \sum_i \lambda_i^*$ can be made arbitrarily close to 0 by increasing the backbone capacity (MLPs as universal function approximators). But practically for fixed n , as enlarging the backbone capacity (i.e., large G , L , and M) also increases ϵ_n and r_n , it is important to choose the backbone network architecture properly. Note also that our assumption on the variational density family for $q_i(\theta_i)$ is easily met; for instance, the families of the mixtures of Gaussians adopted in NIW (Sec. 4.1) and mixture models (Sec. 4.2) obviously subsume a single Gaussian family.*

Proof [Proof of Theorem 2] We first aim to link the variational ELBO objective function to the Hellinger distance via Donsker-Varadhan’s (DV) theorem (Boucheron et al., 2013), motivated from (Bai et al., 2020; Zhang et al., 2022). The DV theorem allows us to express the expectation of any exponential function variationally using the KL divergence. More specifically, the following holds for any distributions p, q and any (bounded) function $h(z)$:

$$\log \mathbb{E}_{p(z)} [e^{h(z)}] = \max_q \left(\mathbb{E}_{q(z)} [h(z)] - \text{KL}(q||p) \right). \quad (139)$$

10. Note that the optimal posterior on ϕ (i.e., $q^*(\phi)$) does not appear here since it only affects $d^2(P_\theta, P^i)$ implicitly. See our detailed analysis/proof provided below.

Here we set $p(z) := p(\theta_i|\phi)$, $q(z) := q_i(\theta_i)$, $h(z) := \log \eta_i(\theta_i)$, where

$$\eta_i(\theta_i) := \exp(l_n(P_{\theta_i}(D_i), P^i(D_i)) + nd^2(P_{\theta_i}, P^i)) \quad \text{and} \quad (140)$$

$$l_n(P_{\theta_i}(D_i), P^i(D_i)) := \log \frac{P_{\theta_i}(D_i)}{P^i(D_i)}, \quad (141)$$

and we have the following inequality that holds for any ϕ :

$$\log \mathbb{E}_{p(\theta_i|\phi)}[\eta_i(\theta_i)] \geq \mathbb{E}_{q_i(\theta_i)}[\log \eta_i(\theta_i)] - \text{KL}(q_i(\theta_i)||p(\theta_i|\phi)). \quad (142)$$

By taking the expectation with respect to $q(\phi)$ and rearranging terms, we have

$$n \cdot \mathbb{E}_{q_i(\theta_i)}[d^2(P_{\theta_i}, P^i)] \leq \mathbb{E}_{q_i(\theta_i)}[-l_n(P_{\theta_i}(D_i), P^i(D_i))] + \mathbb{E}_{q(\phi)}[\text{KL}(q_i(\theta_i)||p(\theta_i|\phi))] + \mathbb{E}_{q(\phi)}[\log \mathbb{E}_{p(\theta_i|\phi)}[\eta_i(\theta_i)]]. \quad (143)$$

For the last term of the right hand side, we use the bound $\mathbb{E}_{s(\theta)}[\eta(\theta)] \leq e^{Cn\epsilon_n^2}$ from the regression theorem (Pati et al., 2018), which holds for any distribution $s(\theta)$ with high probability. The details and proof of this bound can be found in the proof of Theorem 3.1 in (Pati et al., 2018). Applying this bound and telescoping over $i = 1, \dots, N$ yields:

$$n \cdot \sum_{i=1}^N \mathbb{E}_{q_i(\theta_i)}[d^2(P_{\theta_i}, P^i)] \leq \sum_{i=1}^N \mathbb{E}_{q_i(\theta_i)}[-l_n(P_{\theta_i}(D_i), P^i(D_i))] + \sum_{i=1}^N \mathbb{E}_{q(\phi)}[\text{KL}(q_i(\theta_i)||p(\theta_i|\phi))] + N C n \epsilon_n^2. \quad (144)$$

We add $\text{KL}(q(\phi)||p(\phi))$ to the right hand side, which retains the inequality since KL divergence is nonnegative. Then we have the following result that holds for *any* q with high probability,

$$\frac{n}{N} \cdot \sum_{i=1}^N \mathbb{E}_{q_i(\theta_i)}[d^2(P_{\theta_i}, P^i)] \leq L_n(q(\phi), \{q_i(\theta_i)\}_{i=1}^N) + C n \epsilon_n^2, \quad (145)$$

where $L_n(q(\phi), \{q_i(\theta_i)\}_{i=1}^N)$ equals:

$$\frac{1}{N} \left\{ \sum_{i=1}^N \left(\mathbb{E}_{q_i(\theta_i)}[-l_n(P_{\theta_i}(D_i), P^i(D_i))] + \mathbb{E}_{q(\phi)}[\text{KL}(q_i(\theta_i)||p(\theta_i|\phi))] \right) + \text{KL}(q(\phi)||p(\phi)) \right\}, \quad (146)$$

which exactly coincides with our FL-ELBO objective (3) up to constant, by the factor of $1/N$.

Next, we define $\tilde{q}_i(\theta_i)$ and $\tilde{q}(\phi)$ as follows:

$$\tilde{q}_i(\theta_i) = \mathcal{N}(\theta_i; \theta_i^*, \sigma_n^2 I) \quad \text{with} \quad \theta_i^* = \arg \min_{\theta \in \Theta} \|f_\theta - f^i\|_\infty^2, \quad \sigma_n^2 = \frac{G}{8n} A, \quad \text{where} \quad (147)$$

$$A^{-1} = \log(3S_x M) \cdot (2BM)^{2(L+1)} \cdot \left(\left(S_x + 1 + \frac{1}{BM-1} \right)^2 + \frac{1}{(2BM)^2 - 1} + \frac{2}{(2BM-1)^2} \right),$$

$$\tilde{q}(\phi) = \arg \min_{q(\phi)} \sum_{i=1}^N \mathbb{E}_{q(\phi)}[\text{KL}(\tilde{q}_i(\theta_i)||p(\theta_i|\phi))] + \text{KL}(q(\phi)||p(\phi)). \quad (148)$$

Since $(\{q_i^*(\theta_i)\}_{i=1}^N, q^*(\phi))$ is the optimal solution of the FL-ELBO optimisation problem, it is obvious that $L_n(q^*(\phi), \{q_i^*(\theta_i)\}_{i=1}^N) \leq L_n(\tilde{q}(\phi), \{\tilde{q}_i(\theta_i)\}_{i=1}^N)$ if the variational density family for $q_i(\theta_i)$ is rich enough to subsume Gaussian. Now we look closely at $L_n(\tilde{q}(\phi), \{\tilde{q}_i(\theta_i)\}_{i=1}^N)$, and we note that the last two terms as per (146) are constant (i.e., not a function of data size n). That is,

$$\frac{1}{N} \left\{ \sum_{i=1}^N \mathbb{E}_{\tilde{q}(\phi)} [\text{KL}(\tilde{q}_i(\theta_i) \| p(\theta_i | \phi))] + \text{KL}(\tilde{q}(\phi) \| p(\phi)) \right\} = \tilde{C}, \quad (149)$$

for some constant \tilde{C} . Then we can write

$$L_n(\tilde{q}(\phi), \{\tilde{q}_i(\theta_i)\}_{i=1}^N) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\tilde{q}_i(\theta_i)} [-l_n(P_{\theta_i}(D_i), P^i(D_i))] + \tilde{C}. \quad (150)$$

We bound the expected $-l_n$ term in (150) using Lemma 4 below¹¹, which states that with high probability,

$$\mathbb{E}_{\tilde{q}_i(\theta_i)} [-l_n(P_{\theta_i}(D_i), P^i(D_i))] \leq C'n(r_n + \lambda_i^*), \quad (151)$$

for some constant $C' > 0$. Plugging this bound into (150), we have the following derivation where we start from (145) with $q(\phi) = q^*(\phi)$ and $q_i(\theta_i) = q_i^*(\theta_i)$:

$$\frac{n}{N} \cdot \sum_{i=1}^N \mathbb{E}_{q_i^*(\theta_i)} [d^2(P_{\theta_i}, P^i)] \leq L_n(q^*(\phi), \{q_i^*(\theta_i)\}_{i=1}^N) + Cn\epsilon_n^2 \quad (152)$$

$$\leq L_n(\tilde{q}(\phi), \{\tilde{q}_i(\theta_i)\}_{i=1}^N) + Cn\epsilon_n^2 \quad (153)$$

$$\leq \tilde{C} + C'n \left(r_n + \frac{1}{N} \sum_{i=1}^N \lambda_i^* \right) + Cn\epsilon_n^2. \quad (154)$$

By dividing both sides by n , we complete the proof. ■

Lemma 4 (From the proof of Lemma 4.1 in (Bai et al., 2020)). *For $\tilde{q}_i(\theta_i)$ defined¹² as in (147) and r_n, λ_i^* defined as in Theorem 2, the inequality (151) holds with high probability.*

-
11. Although Lemma 4 can be found in the proof of Lemma 4.1 of (Bai et al., 2020), we state this lemma more clearly with separate proof for self-containment.
 12. In (Bai et al., 2020), they defined $\tilde{q}_i(\theta_i)$ as a spike-and-slab model to deal with sparsity. Essentially it is a mixture of two components, selecting $\mathcal{N}(\theta_i; \theta_i^*, \sigma_n^2 I)$ when θ_i^* entries are non-zero and selecting the delta function at 0 when θ_i^* entries equal zero. Without loss of generality (or under mild numerical approximation), we can assume all entries of θ_i^* are non-zero, which makes $\tilde{q}_i(\theta_i)$ Gaussian equal to $\mathcal{N}(\theta_i; \theta_i^*, \sigma_n^2 I)$ as in (147).

Proof [Proof of Lemma 4] From our regression model assumption (133) and (135),

$$\mathbb{E}_{\tilde{q}_i(\theta_i)}[-l_n(P_{\theta_i}(D_i), P^i(D_i))] = \mathbb{E}_{\tilde{q}_i(\theta_i)}[\log P^i(D_i) - \log P_{\theta_i}(D_i)] \quad (155)$$

$$= \frac{1}{2\sigma_\epsilon^2} \left(\mathbb{E}_{\tilde{q}_i(\theta_i)} \|Y^i - f_{\theta_i}(X^i)\|_2^2 - \mathbb{E}_{\tilde{q}_i(\theta_i)} \|Y^i - f^i(X^i)\|_2^2 \right) \quad (156)$$

$$= \frac{1}{2\sigma_\epsilon^2} \left(\underbrace{\mathbb{E}_{\tilde{q}_i(\theta_i)} \|f_{\theta_i}(X^i) - f^i(X^i)\|_2^2}_{\triangleq R_1} + 2 \cdot \underbrace{\mathbb{E}_{\tilde{q}_i(\theta_i)} \langle Y^i - f^i(X^i), f^i(X^i) - f_{\theta_i}(X^i) \rangle}_{\triangleq R_2} \right). \quad (157)$$

We first work on R_1 . Since

$$\|f_{\theta_i}(X^i) - f^i(X^i)\|_2^2 \leq n \cdot \|f_{\theta_i} - f^i\|_\infty^2 \leq 2n \left(\|f_{\theta_i} - f_{\theta_i^*}\|_\infty^2 + \|f_{\theta_i^*} - f^i\|_\infty^2 \right), \quad (158)$$

we have

$$R_1 = 2n \cdot \mathbb{E}_{\tilde{q}_i(\theta_i)} \|f_{\theta_i} - f_{\theta_i^*}\|_\infty^2 + 2n \cdot \|f_{\theta_i^*} - f^i\|_\infty^2 \leq 2n(r_n + \lambda_i^*). \quad (159)$$

where in (159), we use the definition of λ_i^* and the fact $\mathbb{E}_{\tilde{q}_i(\theta_i)} \|f_{\theta_i} - f_{\theta_i^*}\|_\infty^2 \leq r_n$ from Appendix G in (Chérif-Abdellatif, 2020).

Next, we bound R_2 . Since $(Y^i - f^i(X^i)) \sim \mathcal{N}(0, \sigma_\epsilon^2 I)$ and independent of θ_i , we can let $\epsilon := Y^i - f^i(X^i)$ for $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2 I)$. Then

$$R_2 = \epsilon^\top \cdot \mathbb{E}_{\tilde{q}_i(\theta_i)} [f^i(X^i) - f_{\theta_i}(X^i)] \sim \mathcal{N}(0, c_f \sigma_\epsilon^2), \quad (160)$$

where $c_f = \|\mathbb{E}_{\tilde{q}_i(\theta_i)} [f^i(X^i) - f_{\theta_i}(X^i)]\|_2^2$. Applying Jensen's inequality on the convexity of $\|\cdot\|_2^2$, $c_f \leq \mathbb{E}_{\tilde{q}_i(\theta_i)} \|f^i(X^i) - f_{\theta_i}(X^i)\|_2^2 = R_1$. Due to the property of Gaussian, there exists some constant C'_0 such that $R_2 \leq C'_0 \cdot c_f \leq C'_0 \cdot R_1$ with high probability. Plugging these bounds on R_1 and R_2 back to (157) leads to:

$$\mathbb{E}_{\tilde{q}_i(\theta_i)}[-l_n(P_{\theta_i}(D_i), P^i(D_i))] = \frac{1}{2\sigma_\epsilon^2} (R_1 + 2R_2) \leq \frac{1 + 2C'_0}{2\sigma_\epsilon^2} R_1 \leq \frac{1 + 2C'_0}{\sigma_\epsilon^2} n(r_n + \lambda_i^*). \quad (161)$$

Letting $C' := \frac{1+2C'_0}{\sigma_\epsilon^2}$ (constant) completes the proof. \blacksquare

References

- Durmus Alp Emre Acar, Yue Zhao, Ramon Matas Navarro, Matthew Mattina, Paul N Whatmough, and Venkatesh Saligrama. Federated learning based on dynamic regularization. *International Conference on Learning Representations*, 2021.
- Idan Achituve, Aviv Navon, Yochai Yemini, Gal Chechik, and Ethan Fetaya. GP-Tree: A Gaussian Process Classifier for Few-Shot Incremental Learning. In *International Conference on Machine Learning*, 2021a.
- Idan Achituve, Aviv Shamsian, Aviv Navon, Gal Chechik, and Ethan Fetaya. Personalized federated learning with Gaussian processes. In *Advances in Neural Information Processing Systems*, 2021b.

- Maruan Al-Shedivat, Jennifer Gillenwater, Eric King, and Afshin Rostamizadeh. Federated Learning via Posterior Averaging: A New Perspective and Practical Algorithms. In *International Conference on Learning Representations*, 2021.
- Matthew Ashman, Thang D. Bui, Cuong V. Nguyen, Stratis Markou, Adrian Weller, Siddharth Swaroop, and Richard E. Turner. Partitioned Variational Inference: A framework for probabilistic federated learning. *arXiv preprint arXiv:2202.12275*, 2022.
- Jincheng Bai, Qifan Song, and Guang Cheng. Efficient Variational Inference for Sparse Deep Learning with Theoretical Guarantee. In *Advances in Neural Information Processing Systems*, 2020.
- Amir Beck and Luba Tretuashvili. On the convergence of block coordinate descent type methods. *SIAM Journal on Optimization*, 23(4):2037–2060, 2013.
- Dimitri P. Bertsekas. *Nonlinear Programming*. Athena Scientific, 2016.
- Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- Michael Braun and Jon McAuliffe. Variational inference for large-scale models of discrete choice. *arXiv preprint arXiv:0712.2526*, 2008.
- Christopher Briggs, Zhong Fan, and Peter Andras. Federated learning with hierarchical clustering of local updates to improve training on non-iid data. *International Joint Conference on Neural Networks*, 2020.
- Fei Chen, Mi Luo, Zhenhua Dong, Zhenguo Li, and Xiuqiang He. Federated meta-learning with fast convergence and efficient communication. *arXiv preprint arXiv:1802.07876*, 2018.
- Hong-You Chen and Wei-Lun Chao. FedBE: Making Bayesian Model Ensemble Applicable to Federated Learning. In *International Conference on Learning Representations*, 2021.
- Yiqiang Chen, Xin Qin, Jindong Wang, Chaohui Yu, and Wen Gao. Fedhealth: A federated transfer learning framework for wearable healthcare. *IEEE Intelligent Systems*, 35(4):83–93, 2020.
- Badr-Eddine Chérif-Abdellatif. Convergence Rates of Variational Inference in Sparse Deep Learning, 2020. International Conference on Machine Learning.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.
- Canh T Dinh, Nguyen Tran, and Tuan Dung Nguyen. Personalized federated learning with Moreau envelopes. *Advances in Neural Information Processing Systems*, 33, 2020.
- Canh T Dinh, Tung T Vu, Nguyen H Tran, Minh N Dao, and Hongyu Zhang. Fedu: A unified framework for federated multi-task learning with Laplacian regularization. *arXiv preprint arXiv:2102.07148*, 2021.
- Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning: A meta-learning approach. *arXiv preprint arXiv:2002.07948*, 2020.

- Lei Fang, Xiaoli Liu, Xiang Su, Juan Ye, Simon Dobson, Pan Hui, and Sasu Tarkoma. Bayesian Inference Federated Learning for Heart Rate Prediction. *International Conference on Wireless Mobile Communication and Healthcare*, 2020.
- Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *International Conference on Machine Learning*, 2016.
- H. Guo, P. Greengard, H. Wang, A. Gelman, Y. Kim, and E. P. Xing. Federated Learning as Variational Inference: A Scalable Expectation Propagation Approach. *arXiv preprint arXiv:2302.04228*, 2023.
- Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.
- Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3:79–87, 1991.
- M. I. Jordan and R. A. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6(2):181–214, 1994.
- Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J. Reddi, Sebastian U. Stich, and Ananda Theertha Suresh. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. *arXiv preprint arXiv:1910.06378*, 2019.
- D. P. Kingma and M. Welling. Auto-encoding Variational Bayes. In *International Conference on Learning Representations*, 2014.
- Nikita Kotelevskii, Maxime Vono, Eric Moulines, and Alain Durmus. FedPop: A Bayesian Approach for Personalised Federated Learning. In *Advances in Neural Information Processing Systems*, 2022.
- Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated Optimization in Heterogeneous Networks. *arXiv preprint arXiv:1812.06127*, 2018.
- Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated learning through personalization. *International Conference on Machine Learning*, 2021.
- Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the Convergence of FedAvg on Non-IID Data. *arXiv preprint arXiv:1907.02189*, 2019.
- Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of fedavg on non-iid data. In *International Conference on Learning Representations*, 2020.
- Christos Louizos, Matthias Reisser, Joseph Soriaga, and Max Welling. An expectation-maximization perspective on federated learning. *arXiv preprint arXiv:2111.10192*, 2021.
- Wesley Maddox, Timur Garipov, Pavel Izmailov, Dmitry Vetrov, and Andrew Gordon Wilson. A Simple Baseline for Bayesian Uncertainty in Deep Learning. In *Advances in Neural Information Processing Systems*, 2019.
- Yishay Mansour, Mehryar Mohri, Jae Ro, and Ananda Theertha Suresh. Three approaches for personalization with applications to federated learning. *arXiv preprint arXiv:2002.10619*, 2020.

- Othmane Marfoq, Giovanni Neglia, Aurélien Bellet, Laetitia Kamani, and Richard Vidal. Federated Multi-Task Learning under a Mixture of Distributions. In *Advances in Neural Information Processing Systems*, 2021.
- H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data, 2017. AI and Statistics (AISTATS).
- Jaehoon Oh, Sangmook Kim, and Se-Young Yun. FedBABU: Towards Enhanced Representation for Federated Image Classification. In *International Conference on Learning Representations*, 2022.
- Kaan Ozkara, Antonious M. Girgis, Deepesh Data, and Suhas Diggavi. A Statistical Framework for Personalized Federated Learning and Estimation: Theory, Algorithms, and Privacy. In *International Conference on Learning Representations*, 2023.
- D. Pati, A. Bhattacharya, and Y. Yang. On the Statistical Optimality of Variational Bayes, 2018. AI and Statistics (AISTATS).
- Daniel Peterson, Pallika Kanani, and Virendra J. Marathe. Private Federated Learning with Domain Adaptation. *arXiv preprint arXiv:1912.06733*, 2019.
- Lewis J. Rendell, Adam M. Johansen, Anthony Lee, and Nick Whiteley. Global consensus monte carlo. *Journal of Computational and Graphical Statistics*, 30(2), 2021.
- Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet Talwalkar. Federated multi-task learning. *arXiv preprint arXiv:1705.10467*, 2017.
- Benyuan Sun, Hongxing Huo, Yi Yang, and Bo Bai. Partialfed: Cross-domain personalized federated learning via partial initialization. *Advances in Neural Information Processing Systems*, 34, 2021.
- Plassier V., Vono M., Durmus A., and Moulines E. DG-LMC: A turn-key and scalable synchronous distributed MCMC algorithm via Langevin Monte Carlo within Gibbs. In *International Conference on Machine Learning*, 2021.
- Huahua Wang and Arindam Banerjee. Randomized Block Coordinate Descent for Online and Stochastic Optimization. *arXiv preprint arXiv:1407.0107*, 2014.
- Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H. Vincent Poor. Tackling the Objective Inconsistency Problem in Heterogeneous Federated Optimization. *arXiv preprint arXiv:2007.07481*, 2020.
- Stephen J. Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, 2015.
- Hongwei Yang, Hui He, Weizhe Zhang, and Xiaochun Cao. Fedsteg: A federated transfer learning framework for secure image steganalysis. *IEEE Transactions on Network Science and Engineering*, 8(2):1084–1094, 2020.
- Liling Zhang, Xinyu Lei, Yichun Shi, Hongyu Huang, and Chao Chen. Federated Learning with Domain Generalization. *arXiv preprint arXiv:2111.10487*, 2021.
- Xu Zhang, Yinchuan Li, Wenpeng Li, Kaiyang Guo, and Yunfeng Shao. Personalized Federated Learning via Variational Bayesian Inference, 2022. International Conference on Machine Learning.