Keep It Simple: Fault Tolerance Evaluation of Federated Learning with Unreliable Clients

Victoria Huang[*], Shaleeza Sohail[†], Michael Mayo[‡], Tania Lorido Botran[§], Mark Rodrigues[‡], Chris Anderson[‡], Melanie Ooi[‡]

[*]*National Institute of Water and Atmospheric Research*, Wellington, New Zealand
victoria.huang@niwa.co.nz
[†]*University of Newcastle*, New South Wales, Australia
shaleeza.sohail@newcastle.edu.au
[‡]*University of Waikato*, Hamilton, New Zealand
Michael.mayo, Mark.rodrigues, Chris.anderson, Melanie.ooi@waikato.ac.nz
[§]*Roblox*, California, United States
tbotran@roblox.com

Abstract—**Federated learning (FL), as an emerging artificial intelligence (AI) approach, enables decentralized model training across multiple devices without exposing their local training data. FL has been increasingly gaining popularity in both academia and industry. While research works have been proposed to improve the fault tolerance of FL, the real impact of unreliable devices (e.g., dropping out, misconfiguration, poor data quality) in real-world applications is not fully investigated. We carefully chose two representative, real-world classification problems with a limited numbers of clients to better analyze FL fault tolerance. Contrary to the intuition, simple FL algorithms can perform surprisingly well in the presence of unreliable clients.**

Index Terms—**Federated learning, fault tolerance, unreliable clients, robustness, rural environment**

I. Introduction

As an emerging learning paradigm, federated learning (FL) provides a new training method that allows data owners (also called *clients*) to enjoy the performance of a jointly trained model without violating their data privacy. Whereas traditional machine learning (ML) follows a data-to-model approach (i.e. training data from various clients need to be centralized in one node), FL takes an opposite model-to-data approach. In FL, models trained by clients using their local data are shared to a centralized node (called *server*), as opposed to the data. The server updates a global model by aggregating the received local models. The global model will be broadcasted to clients to replace their local models. This means that data never leaves the clients, thus preserving data privacy. Meanwhile, the local sites can take advantage of the improved accuracy resulting from the aggregated model trained across all of the data as opposed to local data only. FL therefore addresses data privacy and data sharing dilemma while potentially offering superior model performance. Given its promising features, FL has attracted much research attention [1].

Although FL has great potential in real-world applications, there have been very few FL applications reported at the

production level, with most of the work being at a proof-of-concept prototype level with synthetic datasets [2]. Meanwhile, implementing FL in a rural environment introduces multiple further complications, all of which are currently active targets for research [3]–[5]. The complications can be broadly grouped into the following issues:

- reducing communication overhead (e.g., using model compression [6]).
- improving energy efficiency (e.g., modifying conventional FL algorithms to save 59% energy cost [7]).
- heterogeneity of clients and where to place computation (e.g., some of the model aggregation can be performed at the edge servers to reduce the communication overhead and also the load on the cloud server [8]).
- impact of unreliable clients that drop in and out of the federation [9], [10] or send low-quality data [11], [12].

This last problem (unreliable clients) is largely unaddressed in the literature and posed as an open research challenge [3].

We therefore focus our study in this paper via a systematic evaluation on the robustness of FL to unreliable clients in rural applications. Note that in this paper, we consider different scenarios of unreliable clients related to *infrastructure-level errors* and *ML-specific inconsistencies*:

- Infrastructure-level errors:
  - Clients randomly and completely drop in and out of the federation. This can be caused by unstable power supplies when clients are deployed in isolated areas. Clients may sometimes drop out during the training process and join back later.
  - Clients randomly and partially drop in and out. Compared to completely dropping out, clients may sometimes fail to upload their locally trained models but they are still able to receive the aggregated global model or vice versa. This can happen when the network connection is unreliable.
- ML-specific inconsistencies:

- Client misconfiguration. For example, the hyper-parameters (e.g., learning rate) of one client are misconfigured and different from others.
- Low-quality data. For example, the training data collected by the client may be low-quality due to hardware constraints or mislabelled due to human mistakes.

The ultimate impact of unreliable clients on a particular system will depend on other factors, such as the number of clients or the size of local dataset [13]. There are efforts to establish metrics that quantify the influence of individual clients on the model outcome, but such metric fails to capture the overall reliability of FL algorithms [14]. Moreover, applications reported in the literature [15], [16] usually involve hundreds of clients. It should be noted that in rural settings, there may only be a very limited number of clients involved, making the impact of unreliable clients potentially significant. Thus, we carefully select our target scenarios that involve *a limited number of clients*.

In this paper, we develop a prototype system building on federated Function-as-a-Service using funcX [17], a federated serverless framework. Two real-world applications have been investigated: weeds detection in precision agriculture in Section III and wildlife detection in camera traps in Section IV. In particular, unreliable clients with infrastructure-level errors will be studied in precision agriculture and ML-specific inconsistencies will be investigated in wildlife detection. To provide representative results, each scenario considers a very low number of clients (between 3 to 6). Extensive experiments have been conducted and showed that even a simple FL algorithm is surprisingly robust to unreliable clients. *By analysing FL robustness against unreliable clients, we aim to understand its impact in real-world applications. This will enable us to evaluate the level of sophistication required of FL algorithms for rural deployment.*

II. PROBLEM FORMULATION

The two target scenarios belong to a classification problem with a set of K data sources (or clients), all scattered over different locations. The statistical properties of each data source are not homogeneous, e.g. different class distributions. Each client is in charge of collecting and storing large amounts of local data. The traditional approach [18], named centralized approach in terms of data collection and model training, involves transferring all the data from each of the clients to a single server that will perform the training. This comes with a high network cost and potentially violates data sovereignty [19].

To circumvent these drawbacks, FL proposes splitting the training between clients and the server [20]. The algorithm is fairly simple: it involves an infinite communication loop between the server and the clients. The server trains a generated (shared) model. A set of clients download a copy (weights) of the shared model and perform local updates leveraging private data, based on some optimization algorithm like stochastic gradient descent. Periodically, clients send a summary of the changes made to the local model (weights of the trained neural network) back to the server. Once it has gathered all the data, the server utilizes some aggregation techniques to perform an update to its shared model [21].

Federated Averaging Algorithm

A very popular aggregation technique is called **Federated Averaging** algorithm [16]. In this case, the server computes the average of the updates received from each client. The full pseudocode is depicted in Algorithm 1. At the server, K clients are selected, which are indexed by a variable k. In parallel, all clients update the generic model weights according to the *ClientUpdate* function, which returns the trained weights w back to the server. Finally, the server computes the average of all weights w received from the K clients. The average of the weights is regarded as the new set of weights for the generic model. Despite its clear advantages, Federated Averaging was not designed with specific fault tolerance features that might lead to sub-optimal model performance. Unreliable clients can be classified into two main types: those with infrastructure-level errors (e.g. faulty network), and the second involves ML-specific inconsistencies (flipped labels, misconfigured hyperparameters such as learning rate, other errors in collected data).

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

---

Server executes ():
    initialize w_0
    for *each round* $t = 1, 2, \ldots$ **do**
        $m \leftarrow max(C \cdot K, 1)$
        $S_t \leftarrow$ (random set of m clients)
        for *each client* $k \in S_t$ **in parallel do**
            $w_{t+1}^k \leftarrow$ **ClientUpdate**(k, w_t)
        end
        $w_{t+1} \leftarrow \sum_{k=1}^{K} \frac{n_k}{n} w_{t+1}^k$
    end

Client update (k, w):
    $B \leftarrow$ (split P_k into batches of size B)
    for *each local epoch i from 1 to E* **do**
        for *batch* $b \in B$ **do**
            $w \leftarrow w - \eta \nabla l(w; b)$
        end
    end
    return w to server;

Infrastructure-level errors in Federated Learning

These errors are due to uncontrollable events in the underlying system. During a global update cycle, the server assumes a total of k weight updates, one belonging to each client. However, one or more clients might fail to send weight updates intermittently or permanently, due to a faulty network

connection [10], [22]. Moreover, as stated earlier, each data source might show different statistical properties, such as class imbalances or even missing representation for certain class(es) [23]. A faulty network coupled with heterogeneous data sources missing might lead to bad model performance, a model unable to generalize well to certain types of cases or even recognize certain classes at all.

A real-world classification problem has been selected to assess the robustness of Federated Averaging learning towards such infrastructure-level issues. The target use case focuses on the detection of certain weed species from hyperspectral images. This kind of dataset requires large volumes of storage space, discouraging from transferring the data to a centralized location. It is the ideal candidate for Federated Learning-based training. The proposed set of experiments empirically analyzes the impact of network failures on the model accuracy, considering both permanent and intermittent failures from one or more clients simultaneously.

ML-specific inconsistencies in Federated Learning

Clients with ML-specific inconsistencies disrupt the model training process, compromise its accuracy or even change the model behavior altogether. The central model relies on properly labelled data from each data source. The primary source of irregularities might come from the data itself: Certain clients might provide mislabelled training data or intentionally flipped class labels [24], [25]. As a result, the resulting model will not be able to detect certain class(es), deeming the system unusable.

Federated learning relies on a number of configuration parameters (see Algorithm 1) to tune the ultimate model accuracy. In certain situations, some clients may have wrong values of such parameters with malicious intentions or simply due to a misconfiguration and will degrade the model performance. The learning rate η has been shown to be a parameter of choice [24], [25] and able to affect the performance of both the aggregated, central model and the local ones [26].

In order to evaluate the effects of ML-specific inconsistencies in federated learning, a second real-world application has been utilized, named animal detection from trap cameras. Such trap cameras are placed on different locations and collect bursts of (at least) three images whenever the motion sensor is triggered. The experiments are designed to analyze the separate effects of mislabelled data and incorrect learning rate values, and also explore whether the combination of both kinds (flipped labels and learning rate values) amplifies the negative effect.

As explained above, Federated Averaging Learning was not designed with fault tolerance behavior. Intuitively, we would expect the model performance to degrade as infrastructure-level or ML- specific inconsistencies develop. The experiment set carried out through two real-world applications will prove (empirically) that Federated Average Learning has a bigger than expected tolerance for such behavior.

III. FL IN PRECISION WEEDS DETECTION

We focus on precision weeds detection as our rural agriculture application. This robotic system uses a near-infrared hyperspectral imaging system to guide the operations of a mobile precision pesticide spraying system [27] (see Fig. 1). The use of near-infrared hyperspectral cameras as a replacement for standard RGB cameras is motivated by research suggesting that spectral measurements can more accurately distinguish weeds from produce [27]. However, the more sophisticated cameras also increase power and data requirements, leading to a need to optimize power and energy consumption in order to achieve maximum area coverage. An effective solution should process camera data locally for weed detection, but also communicate a locally processed model to the global database for aggregation. Therefore, FL is a good fit for this application.

Our intent is to investigate the practical aspects of FL for weeds detection where we deploy a Pasture Care Robot with hyperspectral imaging capabilities. Note that in a large pasture area, full network coverage may not be guaranteed. When the robot is operating in the pasture, it is unavoidable that the robot will completely drop out during the FL training due to network connection issues. Moreover, when the robot's power level is low, it is a common engineering practice that the robot will prioritize navigation over data transmission. In other words, the robot may still receive data but not transmit data in order to conserve battery. Motivated by these, in this weeds detection application, we will investigate the unreliable client scenarios where clients completely or partially drop in and out of the federation.

A. Dataset

As summarized in Table I, the dataset we used in our experiments consists of hyperspectral pasture images taken from three different rural sites [27]. The dataset was labelled with four classes including three different species of pastoral weed: *Jacobaea Vulgaris* (common name: *Ragwort*), *Ulex* (common name: *Gorse*), *Rubus* (common name: *Blackberry*), and a background class of grass. These near-infrared hyperspectral images were taken using a Pika hyperspectral camera and normalised against a white reference to convert from radiance to reflectance while also removing instrument-specific variations. Sample points were extracted from the image datacubes, whereby each pixel had 148 dimensions in the infrared spectrum of 900nm to 1700nm. The dataset was partitioned into training (80% of the samples) and testing (20% of the samples) datasets in a randomized class-stratified manner.

Class normalization of dataset was not performed as our evaluation considers each hyperspectral image individually. In a live environment, it is conceivable that a site has insufficient examples of any given class such that normalization would remove large portions of valid data. We consider this data valid as we compare federated and centralized models to local models: models that would use such data exclusively. A fundamental part of the evaluation of FL in this case is the ability to benefit from the knowledge of other sites.
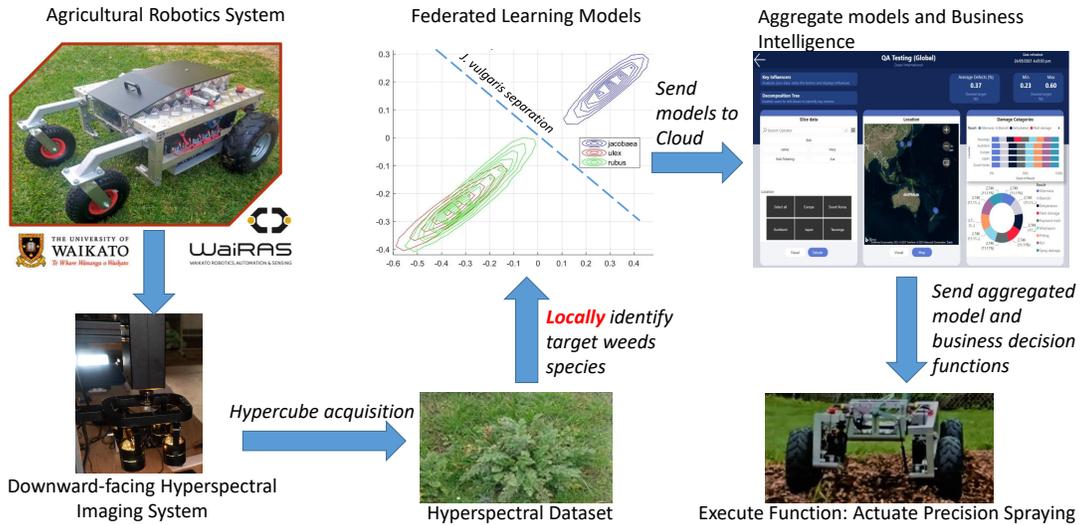
Fig. 1: Hyperspectral Weeds Detection System

TABLE I: Hyperspectral pasture image dataset with imbalanced class distributions and disparate volumes of data among different sites (W: Pastoral weeds, G: Background grass)

Location	Site A	Site B	Site C	Total
#Samples	60,072	30,240	6,232	104,544
#Classes	4 (3W + G)	4 (3W + G)	2 (1W + G)	4 (3W + G)

Note that the heterogeneity in our dataset poses challenges for FL: (1) *Data heterogeneity on class labels*. As shown in Table I, data from the first two sites were balanced across all four of the classes while data from the third consisted of only one weed species and grass, making the distribution of examples between sites quite different and therefore potentially challenging for FL. (2) *Data heterogeneity on data volumes*. Site A provided approximately two times as much data as site B and ten times as much data as site C.

B. Experiment Plan

We evaluate the capabilities of FL in supporting cross-field hyperspectral weeds detection from two aspects:

1) Classification accuracy: To justify the use of FL, we compared FL with a fully localized approach and a fully centralized approach as shown in Fig. 2. The fully localized approach (Fig. 2, green figure) represents a completely decoupled system where each model is trained locally at the three sites and there is no communication at all between sites. The fully centralized approach (Fig. 2, blue figure) represents a completely centralized approach where the model resides on a cloud server and all data must be uploaded to the server before the model is trained. These two approaches represent the current alternative status quo, both of which have disadvantages. For example, the centralized approach can introduce large communication overhead because it requires all data to be uploaded to the cloud. Moreover, the centralized approach cannot be used with sensitive data due to the violation of data sovereignty. Although these issues can be addressed by the
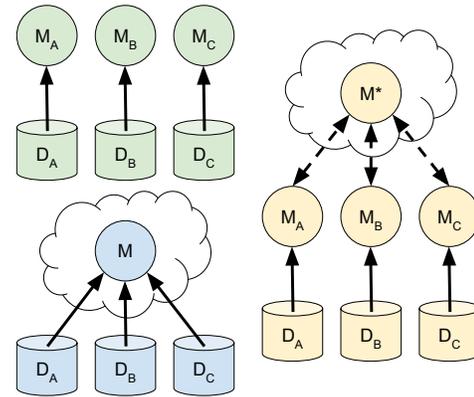
Fig. 2: Approaches to ML in rural AI. Green illustrates a typical fully localised system where data is not shared between sites; each model is locally relevant only with no communication between sites. Blue illustrates a fully centralised system where all data is uploaded to the cloud and a single centralised model is trained; while potentially the most accurate approach, this cannot be used with sensitive data or where the amount of data to be transferred is too large for the available bandwidth. Yellow shows a federated approach where data is kept local while models are sent to the parameter server and used to compute an aggregate model which in turn is shared back to the local sites; this allows learning across sites while not requiring data to be shared.

localized approach, a model trained using local data only may not perform well, especially with a heterogeneous dataset.

On the other hand, in FL (Fig. 2, yellow figure), data resides locally and only the local models are transferred to the cloud server. The cloud server aggregates the models and then returns them back to the local sites. In this way, information and knowledge are shared between the sites indirectly by way of model sharing, as opposed to directly by way of data

TABLE II: Experiment plan

Objective	Research question	Algorithm	Independent variables	Dependent variables
Classification accuracy	How well does FL perform compared with other ML approaches?	FL, Centralized, Localized	Hyperparameters (e.g., iterations, minibatch size, #epochs/iteration)	
Robustness	What is the impact of clients partially dropping out on FL?	FL	When the clients drop out (e.g., upload, download)	Classification accuracy
	What is the effect of training data volume on FL with unreliable clients?		Which client/site drops out (sites with different data volume)	
	What is the effect of client participation rate on FL?		Client participation rate	

uploading. It can be expected that the centralised approach will be the most accurate due to complete access to the data and that the FL approach will be superior to the local-only approach.

2) Robustness against unreliable clients: In rural agricultural farmlands, edge devices (i.e., clients) may sometimes experience unreliable power supply or network connection, which leads to devices dropping in and out during the training process. As a result, a trained model from a local site may not be uploaded to the cloud server for model aggregation. Alternatively, when the failure happens during the model downloading process, the local model is not updated with the aggregated model. To demonstrate the performance of FL for withstanding node failure, we simulate different scenarios under different client participation rates ranging from 100% to 25%. First of all, we investigate the impact of training data volume under the unreliable client setting on the performance of FL. Second, we evaluate how the failures during the model uploading or downloading process affect the performance of FL.

C. Results and Discussion

1) Comparison of different ML approaches:
The first six columns of Table III give a summary of our hyperparameter optimisation results for the short 200 iteration runs. It can be immediately observed that the local models have significantly lower accuracies, with the model at site A achieving a best-case accuracy of marginally above 70%, and the model at site C achieving, at best, 39.5% accuracy. On the other hand, the centralized model achieves nearly 99% test accuracy.

The table also shows the accuracy of the best federated model. This approach, after 200 iterations, reaches the best performance of just over 88% when the number of epochs per iteration is low ($E = 1$) and the batch size is higher ($B = 50$). Unlike the centralized model, however, the federated approach is relatively sensitive to the choice of hyperparameters: specifically, $E = 1$ and $B = 50$ appear to be good choices and we proceeded with further experiments using these hyperparameters.

The results of the further experiments are also given in Table III (rightmost two columns) and also Fig. 3. This column and the figure show that the federated model significantly improves its accuracy with additional iterations. The best model found with FL achieves nearly 94% test accuracy after 2,000 iterations and exceeds 96% accuracy after 10,000 iterations, both good improvements compared to the 200-iteration case.
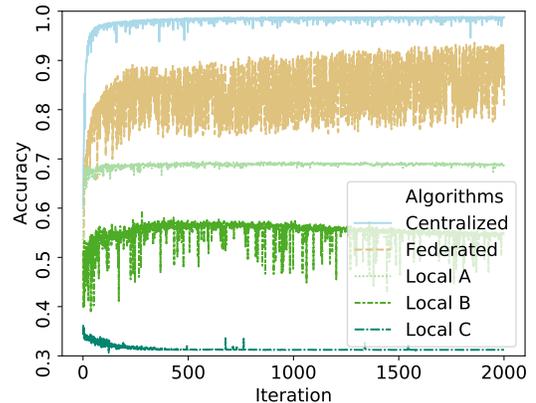
Fig. 3: A comparison between FL, the centralized approach, and the local approach (see Fig. 2) to prove why we should use FL when the centralized approach is not applicable (e.g., high communication cost).

Despite its lack of centralised data, FL can still achieve 96% accuracy as the centralized approach, demonstrating that FL will be a good fit for applications where centralized training is inapplicable. Note that in both cases, the models labelled as "local" are not part of the federation and simulate a situation where models train locally and there is no information sharing between sites at all.

2) FL robustness against client dropout:
Clients completely drop out. Fig. 4 shows the training performance of FL when different clients randomly and completely drop out. For example, in Fig. 4(a), Client A will randomly drop out during the training based on a pre-defined participation rate. From Fig. 4(a), we can see that the training performance is barely affected when Client A only participates 75% of the training iteration. However, larger gaps can be observed when the participation rate drops to 50% or lower. Nevertheless, even with a 75% decrease in participation rate compared to 100% participation, only less than 15% reduction in model accuracy is observed. This result confirms that FL is robust to random client dropout.

The impact of different clients dropout can also be seen by comparing Fig. 4(a), Fig. 4(b), and Fig. 4(c). Due to the heterogeneity in our dataset (as we mentioned in Sec. III-A), the impact of different sites dropping out on FL model accuracy varies. For example, Client A has a more significant impact where almost a 10% accuracy decrease can be spotted with a participation rate decrease from 50% to 25%. On the

TABLE III: Best testing accuracies by hyperparameter combination (columns) and model approach as shown in Fig. 2 (rows). Where $T = 200$, the best accuracies for each approach are bolded. Where $T \geq 2,000$, the accuracy for an approach is bolded only if it exceeds the best accuracy in the preceding columns.

Model Type	(#Epochs/Iteration E, Minibatch Size B, Iterations T)							
	(1,10,200)	(5,10,200)	(20,10,200)	(1,50,200)	(5,50,200)	(20,50,200)	(1,50,2000)	(1,50,10000)
Local (A)	0.693	**0.701**	0.695	0.692	**0.701**	0.699	0.695	0.697
Local (B)	0.571	0.576	0.555	0.569	0.590	0.577	0.592	**0.601**
Local (C)	**0.395**	0.341	0.321	0.382	0.346	0.339	0.363	0.378
Centralized	0.985	**0.988**	**0.988**	0.980	0.987	**0.988**	**0.989**	**0.989**
Federated	0.840	0.809	0.837	**0.883**	0.863	0.861	0.936	**0.963**

(a) Unreliable Client A   (b) Unreliable Client B   (c) Unreliable Client C
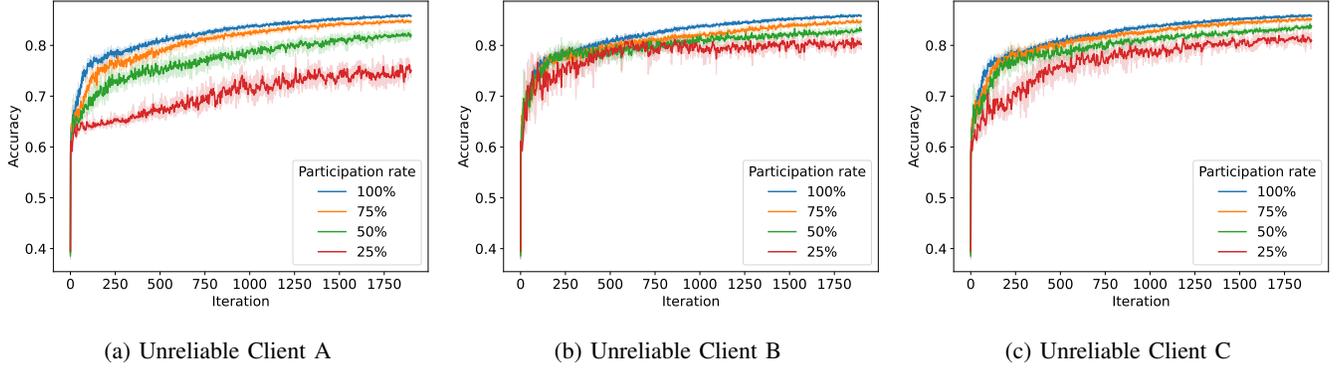
Fig. 4: The training performance of FL with random and complete dropout clients. Each sub-figure shows the FL training performance when a site randomly and completely drops out during FL training process with various probabilities.

(a) Unreliable upload   (b) Unreliable download
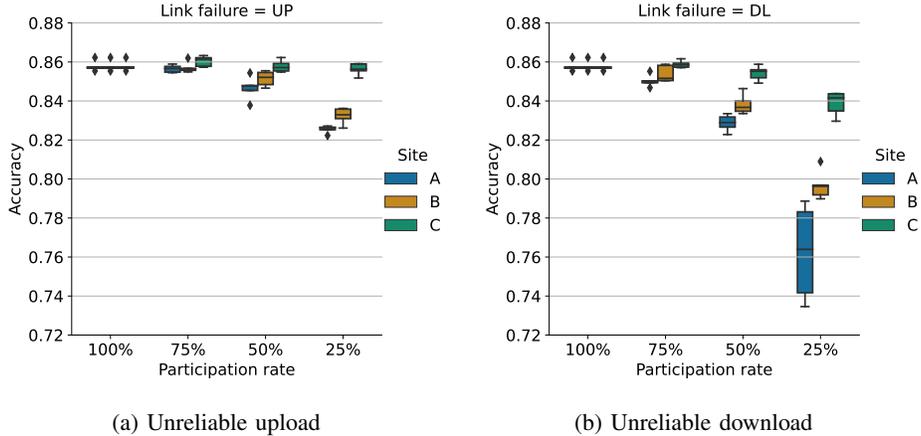
Fig. 5: FL model accuracy with random and partial dropout clients. Note that we select the model after 2,000 training iterations. Each sub-figure demonstrates that under the same link failure, how data heterogeneity affects the FL accuracy. "UP" and "DL" indicate unreliable model upload and download respectively. From both (a) and (b), we can see that site A has the highest impact of the overall accuracy due to its largest data volume and class labels.

other hand, we can only see less than a 5% decrease in accuracy when Client B or C drops out. As we expected, the influence of a client on FL depends on the amount of data and the number of classes it has.

Clients partially drop out. We define partial dropout as clients may randomly fail to upload their local model or download the aggregated model during model training. Fig. 5 and Fig. 6 compared the performance of the models trained at various client participation rates.

Fig. 5(a) shows the impact of clients randomly failing to upload their locally trained models. We can see that there is less than a 4% accuracy decrease when a client does not upload its model. Meanwhile, we also notice that the negative impact is less noticeable if the dropout happens in the client with fewer data, which matches our previous observation. Similar conclusions can also be drawn in Fig. 5(b) when a client fails to download the aggregated models.

In Fig. 6, we compared the different impacts of failing in

(a) Unreliable Client A   (b) Unreliable Client B   (c) Unreliable Client C
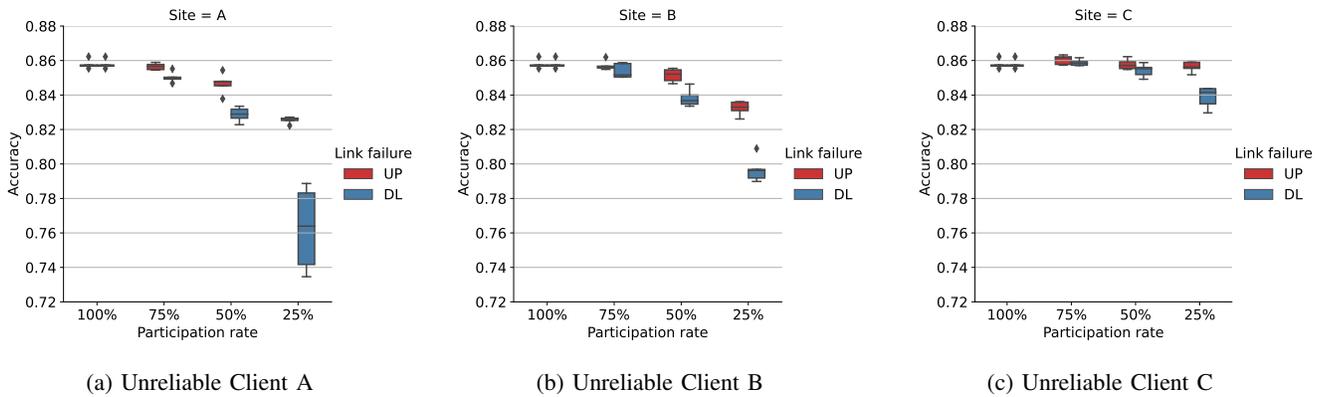
Fig. 6: Performance comparison between different clients and dropout types. Each sub-figure demonstrates that within the same site (i.e., the same training dataset), how different failure types affect the FL accuracy. "UP" and "DL" indicate unreliable model upload and download respectively. For example, from all 3 figures, upload failure barely affects the FL accuracy compared to download failure.

upload or download on the FL model accuracy. For example, we can see from Fig. 6(a) that clients dropping out during model download has a larger impact on the model accuracy. This observation holds for different participation rates as well as different clients.

IV. FL IN CAMERA TRAPS

The current biodiversity crisis requires a global network of remote cameras to monitor and measure essential biodiversity variables using non-invasive remote cameras [28]. A camera trap is a camera with a motion sensor and infrared flash suitable for sampling medium and large-sized birds and mammals with minimal disruption to animal behaviours. There is an increasing trend of using camera traps for monitoring different biodiversity trends, for example, monitoring highway crossing for observing the impact of human activities on wildlife [29], and studying the impact of climate change, population shape, and trophic interactions on elk [30]. Recently, ML has been very successfully used in image-based biodiversity surveys for decision-making purposes by answering critical conservation questions [31]–[33]. However, there are limited coordinated efforts among these camera studies [28]. One limitation is the hesitation in sharing data among parties. Federated learning can play a vital role in this area while maintaining data sovereignty as a large number of participants can create a collaborative ML system incorporating features of all participants' data.

A. Dataset

Wellington camera traps dataset [34] is used to study the robustness of the FL system. This dataset contains 270,450 images, and 90,150 bursts from 187 camera locations in Wellington, New Zealand. The motion sensor cameras recorded a burst of three images when triggered. The images are taken at night and daytime. The dataset is labelled by citizen scientists or professional ecologists from the Victoria University of Wellington, New Zealand. The images are classified into seventeen categories, with a high degree of imbalance across these categories. There is a varying degree of image quality and about 25% of the images are empty, which is considered a false positive.

Several small and camouflaged animals are present in these images, which are very difficult to detect; some example images are shown in Fig. 7. However, the bursts of images provide an opportunity to detect animals by comparing the images in a single burst. [35].

B. Experiment Setup

From the Wellington Camera Trap dataset, we did not consider any image bursts that include less than three images and also did not consider sites that have less than 75 bursts. The image bursts are preprocessed by employing the Mixture of Gaussian (MoG) method and then constructing 3 channel images, as shown in Fig. 8. We have selected different numbers of sites for training, validation, and testing (discussed in the next section). For FL experiments, we have used Flower [36] while using the federated averaging method for client model aggregation. Squeezenet CNN [37] is used with Adam optimiser and starting learning rate of 0.0001. Categorical cross-entropy loss function weighted to account for class imbalance and standard augmentations for training batches [35]. We kept 100 FL iterations with 10 local epochs per iteration and a local batch size of 128.

C. Results and Discussion

The following sets of experiments are conducted to observe the robustness of FL system against ML-specific inconsistencies.

1) FL Robustness against Mislabelled data: ML-specific inconsistency can be due to low-quality data in terms of mislabelled or flipped label records. Several experiments were conducted to study when one or more clients have mislabelled some or all examples in their training data [25]. Such

| (a) True Positive | (b) Possible False Positive | (c) Night Mode | (d) Over Exposed Night Mode |
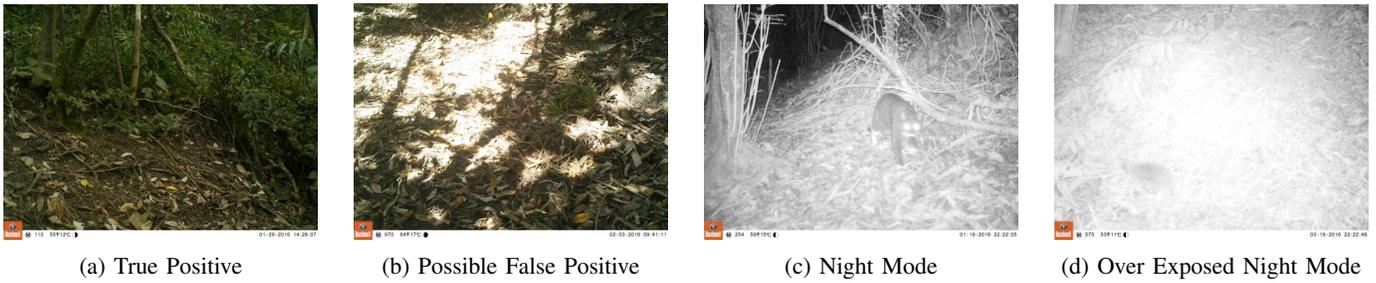
Fig. 7: Examples of single images from Wellington camera trap dataset

Fig. 8: Preprocessing of image bursts
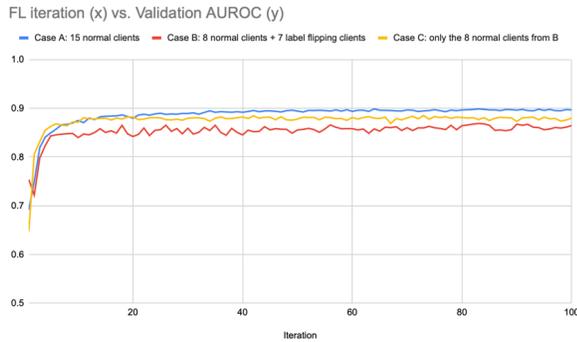
Fig. 9: FL performance under different numbers of mislabelled training data

Fig. 10: FL performance when the learning rate of one client is misconfigured and changed from 0.0001 to 0.01

mislabelled data can affect the global model when the server aggregates the local parameters trained on them. The inherent inaccessibility of the federated server to the clients and clients' data makes it difficult to detect such issues [24].

To evaluate the robustness of FL against such mislabelled clients, we conducted a series of experiments training FL models to compare three scenarios. In the baseline (i.e., Case A), all 15 clients had correctly labelled data. In Case B, 7 clients had mislabelled data while 8 clients had correctly labelled data. Case C involved training the FL model using

data only from the 8 clients with correctly labelled data.

Fig. 9 shows the Area Under the Receiver Operating Characteristics (AUROC) values for all FL iterations for all three cases. Case A shows the AUROC values of around 0.9 after the 20th iteration. Case B has around 47% clients with mislabelled data and shows the AUROC value of 0.86 after the 20th iteration. For case C, where the clients with mislabelled data are removed and the global model considers only 8 normal clients; the AUROC value is around 0.88 after the 20th iteration. We can conclude from the results that the
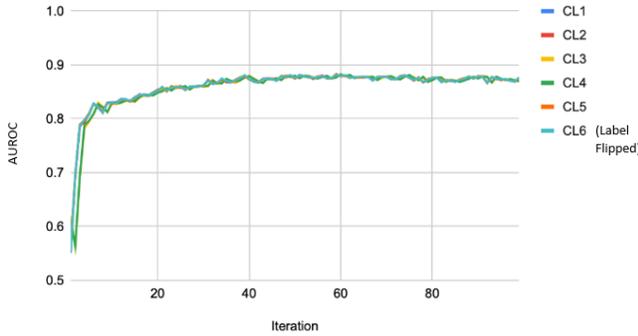
(a) FL performance in the presence of mislabelled data
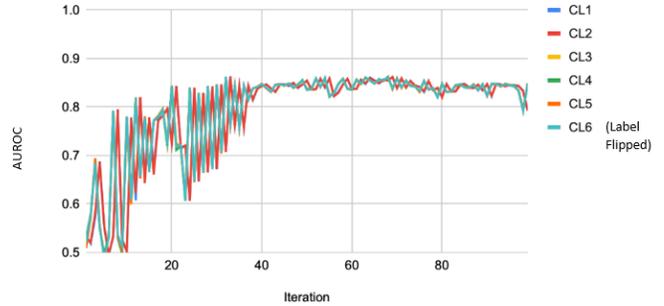
(b) FL performance in the presence of both mislabelled data and misconfigured learning rate of one client (10 times higher)

Fig. 11: The impact of learning rate on each client in the presence of mislabelled data

effect of mislabelled data is there on the performance of the FL model but not very significant, which shows the robustness of FL against such attacks. Some of these findings do not concur with literature [38]. One factor that is really important to consider here is the learning rate; we kept a low learning rate (0.0001) for all our clients, which keeps the changes to the weights very small due to local learning. We have conducted the second set of experiments (discussed next) to study the effect of learning rate in mislabelled data scenarios.

2) FL Robustness against Misconfiguration: One ML-specific inconsistency can be due to the misconfiguration of hyperparameters. One of the hyperparameters that affect the accuracy of the neural networks model is the learning rate, which controls how quickly the model is adapted to the problem [39]. Even for FL systems, the learning rate has shown a significant impact on the accuracy of the models and has been used as one of the factors that can be adapted to improve accuracy [40], fight attacks [41] and flexible client participation [42] among other purposes. We have conducted experiments with six normal clients to study the impact of learning rate on FL systems. We changed the learning rate of only one client from 0.0001 to 0.01; the results are shown in Fig. 10. The results show that for learning rate up to 0.004, the final AUROC values are not affected and is always around 0.88. However, as the learning rate exceeds 0.004, the maximum AUROC cannot go higher than 0.5. This shows that FL shows high robustness against certain misconfigurations (some clients assigned with a higher local learning rate) to a certain extent but as the values become very high, the system cannot perform well.

3) FL Robustness against Mislabelled data and Misconfiguration: The next set of experiments is conducted to study the effect of misconfiguration on the performance of the FL model with different learning rates when some clients have mislabelled data. For these experiments, we have chosen five normal clients and one misconfigured client with mislabelled data. The learning rate for the misconfigured client is changed from 0.0001 to 0.001 to observe the effect on the FL system, as shown in Fig. 11.

The results show that the higher learning rate plays a significant role in the final AUROC as it reduces from 0.88 to 0.85 when the learning rate of the misconfigured client is changed from 0.0001 to 0.001. Also, in early iterations of FL, the impact is more significant and it takes longer for the system to converge in the second case of a higher learning rate. Hence, an important finding is that low local learning rates impact the robustness of the FL systems.

V. RELATED WORK

In this section, we look at some research efforts that tackled unreliable client issues for FL and classify them into two categories: infrastructure-level errors and ML-specific inconsistencies.

Infrastructure-level errors: FL requires sharing of local model parameters between clients and a server over multiple learning rounds until the global model converges. In environments with limited and unreliable network resources, such communication may not be possible. As a result, it can affect the training time and further lead to issues such as unstable and/or slow convergence of the global model. A compensation scheme proposed in [9] suggested that any missing model updates can be replaced by the updates from other clients based on their similarity. Their experiments with the CIFAR-10 dataset show high convergence of the global model under varying successful communication probability for ten clients with a learning rate of 0.001 for every client. A similar approach has been proposed in [10] where it considered local model updates as "friends" when the clients' data distributions are sufficiently similar. To identify "friends", the pairwise similarity score among the local model updates is calculated by the FL server. Whenever any client dropout happens, the FL server replaces the missing model update with its "friends". Experiments were conducted using MNIST and CIFAR-10 datasets. The results showed that model replacement using the friends' model successfully mitigates the effect of client dropout. Recently, a secure and efficient FL (SEFL) framework was proposed in [22]. SEFL considers two servers, one for model aggregation and one for managing cryptography primitives. The clients use the weight pruning technique to

prune the local model update before sending the encrypted updates to the aggregating server. At the server, the encrypted and pruned updates are homomorphically added. In order to decrypt the aggregated value, the other server is contacted. To ensure the security and privacy of the aggregated model values differential privacy approach is used, which is not discussed any further as it is out of the scope of this work. However, a very interesting and related property of the algorithm is that the aggregation server can train an accurate global model even when only 10% of the clients share their local updates with the server.

ML-specific inconsistencies: Data from different clients can be different in terms of quality and label noise. Wrong and/or noisy labels at the clients can negatively affect the global model as the server cannot access local data to filter out the noise. To tackle this problem, different methods have been proposed. For example, a two-level sampling approach from [12] allows a) the server to select better client models and b) clients to select clean local data. The confidence score of data samples is calculated using the global model to calculate the overall confidence of the client. The experiments conducted with the CIFAR-10 dataset with 100 clients containing imbalanced data with varying noise ratios show that the proposed approach outperforms several baseline algorithms. FL under Label Noise (FedLN) [11] estimates per client label noise level and limits the effect of noisy samples on the model's generalisability. At the federated server, Noise-aware Federated Averaging (NA-FedAvg) [11] considers estimated clients' noise levels to perform noise-aware aggregation of the client models. Four different datasets are used to test NA-FedAvg with varying numbers of noisy clients with different noise levels. With all noisy clients, the model performance degraded substantially; however, the presence of only 20% of clean clients enabled the model to perform at 70% for all noise levels. A different approach for reducing the impact of irrelevant or bad quality clients' data on the performance of FL is a distributed selection method [43] that enables clients to choose only a relevant subset from complete data available at the clients. A model requester is proposed for the FL system that provides a benchmarked dataset to the clients to identify a relevant subset of data for a particular FL task.

VI. Conclusions

In this paper, we systematically evaluate the impact of unreliable clients on federated learning (FL) performance in rural environments. We investigated different scenarios of unreliable clients including client dropout, misconfiguration, and low-quality training data. Our experiments are conducted on two applications (weeds detection in precision agriculture and wildlife detection in camera traps) using real-world datasets and a low number of clients (ranging from 3 to 6). Intuitively, scenarios with a reduced number of clients should be more susceptible to performance degradation in case of one or more clients malfunctioning or misconfiguration. Surprisingly, experiments show empirical proof that the federated averaging method is resilient to such unreliable clients,

both to infrastructure-level and ML-specific inconsistencies. First, the weeds detection case study explored the tolerance to the infrastructure-level category of unreliable clients, such as those with faulty network connections. Second, the animal detection case study (via camera traps) showed the robustness of federated learning towards ML-specific deficiencies such as mislabelled data or incorrect parameter values. We have shown that despite its simplicity, federated averaging learning should be the algorithm of choice for distributed training environments with unreliable rural infrastructure and non-technical users.

Our next step involves investigating the applicability of techniques such as Reinforcement Learning as part of the client selection to further improve the performance of Federated Averaging by selecting the best data sources/clients to train the global model. We will further develop platforms and approaches to support the real-world deployment of FL. The platform we used in this paper using Function-as-a-Service provides the first step towards this goal. By introducing higher level systems, programming abstractions, and libraries, we will enable the optimisation workloads and resources of FL in real-world systems.

References

[1] J. Kang, Z. Xiong, D. Niyato, Y. Zou, Y. Zhang, and M. Guizani, "Reliable federated learning for mobile networks," *IEEE Wireless Communications*, vol. 27, no. 2, pp. 72–80, 2020.

[2] S. K. Lo, Q. Lu, C. Wang, H.-Y. Paik, and L. Zhu, "A systematic literature review on federated machine learning: From a software engineering perspective," *ACM Computing Surveys (CSUR)*, vol. 54, no. 5, pp. 1–39, 2021.

[3] W. Y. B. Lim, N. C. Luong, D. T. Hoang, Y. Jiao, Y.-C. Liang, Q. Yang, D. Niyato, and C. Miao, "Federated learning in mobile edge networks: A comprehensive survey," *IEEE Communications Surveys Tutorials*, vol. 22, no. 3, pp. 2031–2063, 2020.

[4] P. Patros, M. Ooi, V. Huang, M. Mayo, C. Anderson, S. Burroughs, M. Baughman, O. Almurshed, O. Rana, R. Chard *et al.*, "Rural ai: Serverless-powered federated learning for remote applications," *IEEE Internet Computing*, 2022.

[5] O. Almurshed, P. Patros, V. Huang, M. Mayo, M. Ooi, R. Chard, K. Chard, O. Rana, H. Nagra, M. Baughman *et al.*, "Adaptive edge-cloud environments for rural ai," in *2022 IEEE International Conference on Services Computing (SCC)*. IEEE, 2022, pp. 74–83.

[6] Y. Xu, Y. Liao, H. Xu, Z. Ma, L. Wang, and J. Liu, "Adaptive control of local updating and model compression for efficient federated learning," *IEEE Transactions on Mobile Computing*, 2022.

[7] Z. Yang, M. Chen, W. Saad, C. S. Hong, and M. Shikh-Bahaei, "Energy efficient federated learning over wireless communication networks," *CoRR*, vol. abs/1911.02417, 2019. [Online]. Available: http://arxiv.org/abs/1911.02417

[8] L. Liu, J. Zhang, S. Song, and K. B. Letaief, "Client-edge-cloud hierarchical federated learning," in *ICC 2020-2020 IEEE International Conference on Communications (ICC)*. IEEE, 2020, pp. 1–6.

[9] Y. Mao, Z. Zhao, M. Yang, L. Liang, Y. Liu, W. Ding, T. Lan, and X.-P. Zhang, "Safari: Sparsity enabled federated learning with limited and unreliable communications," *arXiv preprint arXiv:2204.02321*, 2022.

[10] H. Wang and J. Xu, "Friends to help: Saving federated learning from client dropout," *arXiv preprint arXiv:2205.13222*, 2022.

[11] V. Tsouvalas, A. Saeed, T. Ozcelebi, and N. Meratnia, "Federated learning with noisy labels," *arXiv preprint arXiv:2208.09378*, 2022.

[12] Z. Wang, T. Zhou, G. Long, B. Han, and J. Jiang, "Fednoil: A simple two-level sampling method for federated learning with noisy labels," *arXiv preprint arXiv:2205.10110*, 2022.

[13] M. Kamp, J. Fischer, and J. Vreeken, "Federated learning from small datasets," *arXiv preprint arXiv:2110.03469*, 2021.

[14] Y. Xue, C. Niu, Z. Zheng, S. Tang, C. Lyu, F. Wu, and G. Chen, "Toward understanding the influence of individual clients in federated learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 12, 2021, pp. 10 560–10 567.

[15] A. Nilsson, S. Smith, G. Ulm, E. Gustavsson, and M. Jirstrand, "A performance evaluation of federated learning algorithms," in *Proceedings of the second workshop on distributed infrastructures for deep learning*, 2018, pp. 1–8.

[16] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Singh and J. Zhu, Eds., vol. 54. Fort Lauderdale, FL, USA: PMLR, 20–22 Apr 2017, pp. 1273–1282. [Online]. Available: http://proceedings.mlr.press/v54/mcmahan17a.html

[17] R. Chard, Y. Babuji, Z. Li, T. Skluzacek, A. Woodard, B. Blaiszik, I. Foster, and K. Chard, "FuncX: A federated function serving fabric for science," in *29th International Symposium on High-Performance Parallel and Distributed Computing*, 2020, pp. 65–76.

[18] P. Wang, E. Fan, and P. Wang, "Comparative analysis of image classification algorithms based on traditional machine learning and deep learning," *Pattern Recognition Letters*, vol. 141, pp. 61–67, 2021.

[19] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning," in *2019 IEEE symposium on security and privacy (SP)*. IEEE, 2019, pp. 739–753.

[20] Q. Yang, Y. Liu, Y. Cheng, Y. Kang, T. Chen, and H. Yu, "Federated learning," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 13, no. 3, pp. 1–207, 2019.

[21] E. Sannara, F. Portet, P. Lalanda, and V. German, "A federated learning aggregation algorithm for pervasive computing: Evaluation and comparison," in *2021 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 2021, pp. 1–10.

[22] J. Deng, C. Wang, X. Meng, Y. Wang, J. Li, S. Lin, S. Han, F. Miao, S. Rajasekaran, and C. Ding, "A secure and efficient federated learning framework for nlp," *arXiv preprint arXiv:2201.11934*, 2022.

[23] L. Wang, S. Xu, X. Wang, and Q. Zhu, "Addressing class imbalance in federated learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 11, 2021, pp. 10 165–10 173.

[24] Z. Lv, H. Cao, F. Zhang, Y. Ren, B. Wang, C. Chen, N. Li, H. Chang, and W. Wang, "Awfc: Preventing label flipping attacks towards federated learning for intelligent iot," *The Computer Journal*, vol. 65, no. 11, pp. 2849–2859, 2022.

[25] N. M. Jebreel, J. Domingo-Ferrer, D. Sánchez, and A. Blanco-Justicia, "Defending against the label-flipping attack in federated learning," *arXiv preprint arXiv:2207.01982*, 2022.

[26] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin, and V. Shmatikov, "How to backdoor federated learning," in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, S. Chiappa and R. Calandra, Eds., vol. 108. PMLR, 26–28 Aug 2020, pp. 2938–2948. [Online]. Available: https://proceedings.mlr.press/v108/bagdasaryan20a.html

[27] W. Holmes, M. P.-L. Ooi, M. Look, Y. C. Kuang, R. Simpkin, D. Blanchon, and S. Demidenko, "Proximal near-infrared spectral reflectance characterisation of weeds species in New Zealand pasture," in *IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*. IEEE, 2019, pp. 1–6.

[28] R. Steenweg, M. Hebblewhite, R. Kays, J. Ahumada, J. T. Fisher, C. Burton, S. E. Townsend, C. Carbone, J. M. Rowcliffe, J. Whittington *et al.*, "Scaling-up camera traps: Monitoring the planet's biodiversity with networks of remote sensors," *Frontiers in Ecology and the Environment*, vol. 15, no. 1, pp. 26–34, 2017.

[29] M. Barrueto, A. T. Ford, and A. P. Clevenger, "Anthropogenic effects on activity patterns of wildlife at crossing structures," *Ecosphere*, vol. 5, no. 3, pp. 1–19, 2014.

[30] J. F. Brodie, E. Post, J. Berger, and F. Watson, "Trophic interactions and dynamic herbivore responses to snowpack," *Climate Change Responses*, vol. 1, no. 1, pp. 1–8, 2014.

[31] I. Zualkernan, S. Dhou, J. Judas, A. R. Sajun, B. R. Gomez, and L. A. Hussain, "An iot system using deep learning to classify camera trap images on the edge," *Computers*, vol. 11, no. 1, p. 13, 2022.

[32] O. Pantazis, G. J. Brostow, K. E. Jones, and O. Mac Aodha, "Focus on the positives: Self-supervised learning for biodiversity monitoring," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 583–10 592.

[33] M. S. Norouzzadeh, D. Morris, S. Beery, N. Joshi, N. Jojic, and J. Clune, "A deep active learning system for species identification and counting in camera trap images," *Methods in ecology and evolution*, vol. 12, no. 1, pp. 150–161, 2021.

[34] V. Anton, S. Hartley, A. Geldenhuis, and H. U. Wittmer, "Monitoring the mammalian fauna of urban areas using remote cameras and citizen science," *Journal of Urban Ecology*, vol. 4, no. 1, p. juy002, 2018.

[35] B. M. Shashidhara, D. Mehta, Y. Kale, D. Morris, and M. Hazen, "Sequence information channel concatenation for improving camera trap image burst classification," *arXiv preprint arXiv:2005.00116*, 2020.

[36] D. J. Beutel, T. Topal, A. Mathur, X. Qiu, J. Fernandez-Marques, Y. Gao, L. Sani, K. H. Li, T. Parcollet, P. P. B. de Gusmão *et al.*, "Flower: A friendly federated learning framework," 2022.

[37] V. Pothos, D. Kastaniotis, I. Theodorakopoulos, and N. Fragoulis, "A fast, embedded implementation of a convolutional neural network for image recognition," 2016.

[38] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," *Knowledge-Based Systems*, vol. 216, p. 106775, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0950705121000381

[39] T. Takase, S. Oyama, and M. Kurihara, "Effective neural network training with adaptive learning rate based on training loss," *Neural Networks*, vol. 101, pp. 68–78, 2018.

[40] C. Xu, S. Liu, Z. Yang, Y. Huang, and K.-K. Wong, "Learning rate optimization for federated learning exploiting over-the-air computation," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 12, pp. 3742–3756, 2021.

[41] M. S. Ozdayi, M. Kantarcioglu, and Y. R. Gel, "Defending against backdoors in federated learning with robust learning rate," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 10, 2021, pp. 9268–9276.

[42] H. Yang, X. Zhang, P. Khanduri, and J. Liu, "Anarchic federated learning," in *International Conference on Machine Learning*. PMLR, 2022, pp. 25 331–25 363.

[43] T. Tuor, S. Wang, B. J. Ko, C. Liu, and K. K. Leung, "Overcoming noisy and irrelevant data in federated learning," in *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021, pp. 5020–5027.