

**Federated learning framework for collaborative remaining useful life prognostics
An aircraft engine case study**

Landau, Diogo; de Pater, Ingeborg; Mitici, Mihaela; Saurabh, Nishant

DOI

[10.1016/j.future.2025.107945](https://doi.org/10.1016/j.future.2025.107945)

Publication date

2026

Document Version

Final published version

Published in

Future Generation Computer Systems

Citation (APA)

Landau, D., de Pater, I., Mitici, M., & Saurabh, N. (2026). Federated learning framework for collaborative remaining useful life prognostics: An aircraft engine case study. *Future Generation Computer Systems*, 174, Article 107945. <https://doi.org/10.1016/j.future.2025.107945>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Federated learning framework for collaborative remaining useful life prognostics: An aircraft engine case study

Diogo Landau^a ,^{*} Ingeborg de Pater^b ,¹ Mihaela Mitici^a, Nishant Saurabh^a 

^a Department of Information and Computing Sciences, Utrecht University, Heidelberglaan 8, 3584 CS Utrecht, The Netherlands

^b Faculty of Aerospace Engineering, Delft University of Technology, Kluyverweg 1, HS 2926 Delft, The Netherlands

ARTICLE INFO

Keywords:

Federated and collaborative learning
Remaining useful life prognostics
Robust parameter aggregation method
Distributed data
N-CMAPSS turbofan engines

ABSTRACT

Complex systems such as aircraft engines are continuously monitored by sensors. In predictive aircraft maintenance, the collected sensor measurements are used to estimate the health condition and the Remaining Useful Life (RUL) of such systems. However, a major challenge when developing prognostics is the limited number of run-to-failure data samples. This challenge could be overcome if multiple airlines would share their run-to-failure data samples such that sufficient learning can be achieved. Due to privacy concerns, however, airlines are reluctant to share their data in a centralized setting. In this paper, a collaborative federated learning framework is therefore developed instead. Here, several airlines cooperate to train a collective RUL prognostic machine learning model, without the need to centrally share their data. For this, a decentralized validation procedure is proposed to validate the prognostics model without sharing any data. Moreover, sensor data is often noisy and of low quality. This paper therefore proposes four novel methods to aggregate the parameters of the global prognostic model. These methods enhance the robustness of the FL framework against noisy data. The proposed framework is illustrated for training a collaborative RUL prognostic model for aircraft engines, using the N-CMAPSS dataset. Here, six airlines are considered, that collaborate in the FL framework to train a collective RUL prognostic model for their aircraft's engines. When comparing the proposed FL framework with the case where each airline independently develops their own prognostic model, the results show that FL leads to more accurate RUL prognostics for five out of the six airlines. Moreover, the novel robust aggregation methods render the FL framework robust to noisy data samples.

1. Introduction

The emergence of distributed computing technologies integrated with machine learning techniques, has gained traction for distributed data processing across various application domains. Predictive maintenance, and in particular the generation of data-driven Remaining Useful Life (RUL) prognostics for collaborative systems, has benefited in recent years from the integration of distributed computing technologies with machine learning models [1,2].

The main challenge when developing RUL prognostics using machine learning is the lack of run-to-failure data samples [1,3]. This is because, although systems are continuously monitored by sensors, they are often maintained/replaced preventively before their actual failure (preventive maintenance). This is, for example, the case for safety-critical aircraft systems such as the engines, where preventive maintenance is a means to limit the risk and consequences of unexpected failures. Most condition-monitoring data samples of such

systems are unlabeled, i.e., the true RUL or health state is unknown. Only in a few cases are the systems monitored until their actual failure, i.e., there are very few run-to-failure data samples or, equivalently, there are very few labeled data samples for which the true RUL is known. This lack of run-to-failure data samples is a challenge when developing RUL prognostics using machine learning models.

One way to overcome this challenge is for clients (airlines), each with their own condition-monitoring data samples, to collaborate in developing a common RUL prognostics model. When clients collaborate and share labeled data samples, valuable data are made available to train machine learning models for RUL prognostics. In the case of prognostics development for aircraft systems, however, the airlines are very hesitant to share their data, or to transmit the data for processing to a centralized server, due to privacy concerns and possible conflicts of interests [4–6]. The aversion to data sharing remains even when the data is anonymized or processed [4].

* Corresponding author.

E-mail address: d.hewittmouracarrapatolandau@uu.nl (D. Landau).

¹ These authors contributed equally to this work.

In this paper, a collaborative federated learning (FL) framework is proposed for developing RUL prognostics for aircraft systems. For this purpose, a common machine learning RUL prognostic model is trained using the data of several airlines. It is assumed that each airline (client) possesses a local compute device. Collaborative FL is initially deployed across these local devices [7], similar to classic FL, such that the clients (airlines) do not have to share or transmit data to a centralized server to train this common model. Firstly, it makes use of the intrinsic value of data samples of multiple airlines, further recognizing that the collective intelligence of multiple airline data samples can enable the development of accurate and robust RUL prognostics. Secondly, it aligns with the imperative to preserve data privacy, an issue of paramount importance for airlines. By taking condition monitoring data processing and prognostics model training to the individual airlines, this approach minimizes data transfer, and empowers airlines to collaboratively generate RUL prognostics, while safeguarding sensitive operational information.

From a methodological perspective, there are two main aspects to consider when developing a collaborative FL framework for RUL prognostics for aircraft systems. Firstly, when clients collaborate, the quality of the data has a direct impact on the final collaborative model. For example, condition-monitoring data samples are often noisy due to interference from the outside environment, data transmission, various data sources and various sensors [8,9]. Such noisy data directly impacts the accuracy of the common model [4,10–12]. Standard aggregation procedures used for classic FL, e.g., FedAvg [7], are unable to prevent clients with noisy data from participating in the development of a common model, leading to a decreased accuracy of the model. To mitigate this, recent studies [13,14] have provided robust algorithms that guarantee model convergence in the presence of noisy contributions. However, these methods assume that the data is independent and identically distributed (i.i.d.) between the different clients, which is rarely the case for airlines. For instance, each airlines flies different routes with their aircraft, which influences the sensor measurements. As such, this paper presents robust parameter aggregation procedures for a collaborative FL-based RUL prognostics that are robust to noisy data samples.

Secondly, a decentralized validation procedure is formalized for validating a RUL prognostic machine learning model, embedded in the proposed FL framework. This decentralized validation procedure is used to train the RUL prognostics models. Generally, classic FL is performed (i) without any validation procedure [15–19], or (ii) by using a global validation set [4,11,20,21]. However, neither of these two cases is applicable in aviation. First, the European Union Aviation Safety Agency (EASA) requires that the machine learning models are validated [22] in order to approve their usage. Without a validation procedure, machine learning models for RUL prognostics of aircraft systems cannot be certified and implemented in practice. Second, the alternative of having a global validation set would require the airlines to share some data with the server, which violates their privacy. As such, assuming the availability of a global validation set for RUL prognostics of aircraft systems, as assumed in [4,11,20], is not realistic. To address this, the proposed decentralized validation procedure enables each client (airline) to validate the global model with their own validation dataset, and share only the loss obtained. With this, a global validation loss is determined. Overall, the proposed decentralized validation procedure supports the certification of machine learning models for RUL prognostics of aircraft systems, without requiring the airlines to share their data.

The main contributions of this paper are as follows:

- A collaborative FL framework is proposed to train a RUL prognostic machine learning model with multiple clients (airlines) without sharing any data. For this, a decentralized validation procedure is formalized to validate the machine learning models, which is embedded in the FL framework. As data privacy preservation is crucial in such settings, this decentralized validation

procedure does not require the clients (airlines) to share any data or labels. Nevertheless, this decentralized validation procedure is generic, and can be employed for other applications of FL.

- Four robust approaches are developed to aggregate the parameters for the global model. This enables the development of accurate RUL prognostics even when some clients (airlines) have noisy data samples.
- The proposed framework is applied to a realistic case study, where 6 airlines collaborate to train a RUL prognostic model for aircraft engines. For this the N-CMAPSS dataset [23] on aircraft engines is used. The proposed FL approach is compared with the case when airlines do not collaborate at all and each airline independently trains its own RUL prognostic models. In this case, the proposed FL approach leads to more accurate RUL prognostics for 5 out of the 6 considered airlines. The complete source code of the proposed FL framework including the reproducibility and data artifact is released opensource in an *online* GitHub repository [24].

The rest of the paper is organized as follows. In Section 3, an overview of the FL approach is provided, which includes the decentralized validation procedure. In Section 4 this FL approach is further extended with noise-robust aggregation procedures for the model parameters. In Section 5 the proposed approach is illustrated in a case study where airlines collaboratively generate RUL prognostics. Numerical results are provided in Section 6, and conclusions are provided in Section 7.

2. Related work

Recent studies employed FL in a variety of applications, such as healthcare [25], smart cities [26] and the Internet of Things [27,28]. Most studies on FL in the domain of condition monitoring and health management of systems focus on fault diagnosis. For instance, in [4, 15–19,29], the faults of bearings are diagnosed.

The development of a RUL prognostics model with FL is very rare in literature. In [11], a convolutional autoencoder is trained at each client. The clients then send the encoded state of each data sample, together with the corresponding label to the server. This server trains a RUL prognostic model on these encoded states. For sensitive applications, however, the encoded state of the data samples and corresponding labels may already be too valuable to share. Therefore, in this paper, airlines simply share their local models that resulted from training on their local datasets. In [30], FL is used to train a RUL prognostic model for clients with heterogeneous data. Here, the clients are first clustered based on the similarity of their data, and RUL prognostic models are trained based on these clusters. This paper, however, considers a problem where airlines train a common RUL prognostic model through FL. As not many airlines exist, and as each airlines only has very few labeled data samples, this RUL prognostic model is trained with all participating airlines.

2.1. Aggregation methods

Given the increased statistical heterogeneity in FL compared to other distributed learning settings, aggregation methods have been a core concern in the literature. Li et al. [31] modify *FedAvg* to account for statistical and system heterogeneity, while Ji et al. [32] leverage a minimization problem to compute the global model that is closest (in distance) to the submitted local models. However, these procedures do not account for noisy local models that negatively contribute to the global aggregation procedure, a scenario that if unhandled, deters potential participants.

For this reason, several robust aggregation procedures [33–36] have been proposed to mitigate the impact clients with noisy data can have on the performance of the global model. Given a set of models

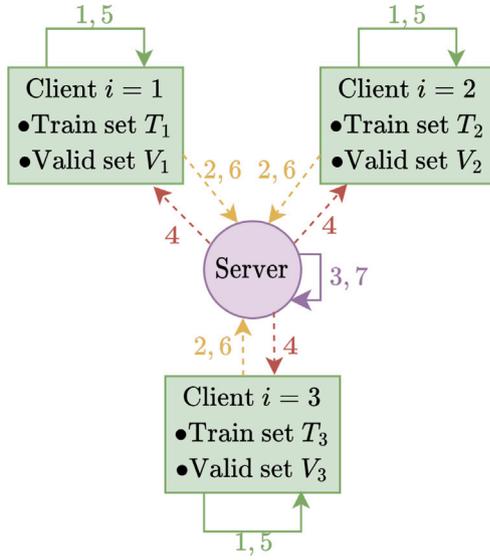


Fig. 1. Schematic overview of FL with the proposed decentralized validation procedure, $N = 3$ clients.

trained by local clients, robust aggregation methods aim to maintain the convergence of the global model regardless of the presence of noise [13,33] or malicious intent [14,37]. Blanchard et al. [14] select a single model as the global model based on a distance metric computed for all local models. Yin et al. [13] instead propose two methods that compute the global model based on the distribution of the parameters over all local models. Nevertheless, as mentioned by the authors and empirically evaluated in [21], these methods assume the data is independent and identically distributed, an assumption that does not hold in the considered scenario where the data is held by different airline companies.

To circumvent this, Fang et al. [21] present an aggregation method that evaluates the performance of each local model on a *global* validation dataset, followed by removing the models that have a negative impact on the global model's performance. Their use of a global validation dataset hinders its use in aviation, wherein multiple airlines would have to share data to construe a representative dataset for the server. For this reason, this paper presents a decentralized validation procedure to inform robust aggregation methods in the context of RUL prognostics.

3. Federated learning with a decentralized validation procedure

In this section, a FL approach with a decentralized validation procedure is proposed, where the clients do not have to share any data or data labels. This approach is motivated by the validation procedure presented in [38–41], where the clients possess a local test set to evaluate the performance of a given model. This approach is applicable to a broad set of machine learning problems (e.g. regression or classification problems).

Let C be a set of clients (e.g., airlines, companies, transport companies), where each client $i \in C$ has its own dataset X_i . This dataset X_i only belongs to client i , and cannot be shared with other clients or with the central server. In total, there are $|C| = N$ clients. For instance, for RUL prognostics, the data set X_i consists of data samples with sensor measurements, where each data sample has as label the actual RUL. Standard FL uses all data samples in the dataset X_i of client $i \in C$ as training set. In contrast, in this paper the data samples at each client $i \in C$ are divided into a training set $T_i \subset X_i$ and a validation set $V_i = X_i \setminus T_i$. Thus, the training-validation split is performed at each client separately.

1. Train global model for one epoch with T_i , obtain W_i
5. Validate global model with V_i , obtain sum of loss $\mathcal{L}^{\text{sum}}(V_i, W_G)$

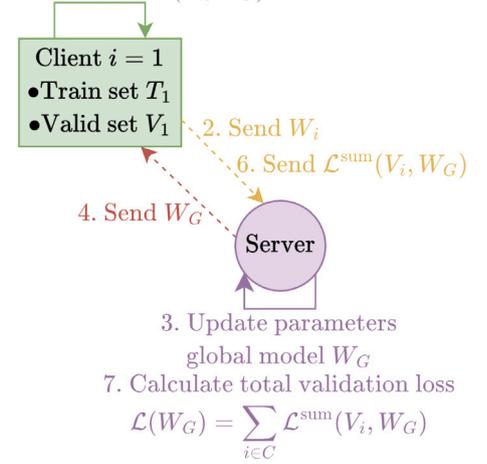


Fig. 2. Zoom in - Schematic overview of all steps of the FL with the proposed decentralized validation procedure, zoomed in on one client and the server.

The clients cooperate and communicate with each other through a central server, without sharing any data. This server contains the global model, which is trained jointly by all clients. In this Section, it is assumed that this global model is a neural network. Let W_G denote the parameters of this global model. It is also assumed that no global validation set is available at this server.

The server first initializes the global model. The parameters W_G are then sent to all clients $i \in C$. After this, the procedure for training the global model for one epoch with FL with the decentralized validation procedure is as follows (see also Figs. 1 and 2):

1. Each client $i \in C$ trains the global model for one epoch using its training set T_i . The resulting neural network is referred to as the client's local model. Let W_i denote the parameters of the local model of client $i \in C$. Note that only a part of the data ($T_i \subset X_i$) at each client $i \in C$ is used to train the model.
2. Each client $i \in C$ sends the parameters W_i to the server.
3. The server aggregates the parameters W_i from all clients $i \in C$, to generate the new global model. Each parameter $w_G \in W_G$ of the global model is a function $f(\cdot)$ of the corresponding parameter $w_i \in W_i$ of the local model of all clients $i \in C$, i.e.,:

$$w_G = f(w_1, w_2, \dots, w_N). \quad (1)$$

The standard aggregation procedure is called federated averaging (FedAvg). In FedAvg, each parameter $w_G \in W_G$ of the global model is simply the weighted average of this parameter $w_i \in W_i$ over the models of all clients $i \in C$ [7]:

$$w_G = f(w_1, w_2, \dots, w_N) = \sum_{i \in C} \frac{|T_i|}{n} w_i, \quad (2)$$

where $|T_i|$ is the number of samples in the training set T_i of client $i \in C$, while n is the total number of samples at all clients, i.e., $n = \sum_{i \in C} |T_i|$.

4. The server sends the updated parameters W_G of the global model to all clients $i \in C$.
5. Each client $i \in C$ validates the global model (with parameters W_G) with its own validation dataset V_i , by computing the sum of the loss $\mathcal{L}^{\text{sum}}(V_i, W_G)$. This is, for instance, the sum of the squared errors for regression problems, or the sum of the cross entropy loss for classification problems.

6. Each client $i \in C$ sends the loss $\mathcal{L}^{\text{sum}}(V_i, W_G)$ back to the server.
7. The performance of the global model is assessed by calculating the total validation loss $\mathcal{L}(W_G)$ from the losses $\mathcal{L}^{\text{sum}}(V_i, W_G)$ of all clients $i \in C$:

$$\mathcal{L}(W_G) = \sum_{i \in C} \mathcal{L}^{\text{sum}}(V_i, W_G). \quad (3)$$

Note that Steps 5, 6 and 7 are not present in a classical FL procedure without any validation.

The model is trained until a certain stopping criterion is fulfilled, for example based on the number of epochs or on the total validation loss. Due to the decentralized validation procedure, after training, the parameters of the global model that result in the lowest validation loss can be selected.

4. Noise-robust federated learning aggregation methods for model parameters

In Section 3, Steps 1–7 illustrate how a global model is generated and validated in a general distributed FL setup. With classical aggregation methods such as FedAvg, Step 3 computes a global model based on a weighted average of the local models provided by each client (Eq. (2)). With this approach, the contribution each model has in the global model is proportional to the size of the dataset used by each client. However, given this collaborative FL setup, it is possible that clients train their models on noisy data due to, for example, faulty sensors. In these conditions, due to the averaging performed by FedAvg, the noise has an adverse effect on the global model’s performance. To mitigate such issues, in this Section, Step 3 is split into Steps 3a and 3b to add an evaluation procedure before computing the global model. The intuition behind this design, is that by evaluating the performance of the models before aggregation, a model’s contribution to the global model can be tuned based on its performance on different datasets available at each client. Specifically, the weights associated with each model are tuned to increase the contribution of the best performing models, and decrease the contribution of the models that do not perform as well.

With this in mind, four parameter aggregation methods are presented to increase the resilience of the global model against noisy data. These aggregation methods are based on the proposed decentralized validation procedure. Here, the local validation sets are used to evaluate the importance of each client’s local model for the aggregation procedure. The evaluation is performed every time a new global model is to be created. For this, step 3 in the FL framework (Section 3) is split in two sub-steps:

- Step 3a: First, it is evaluated how well a local model of client $i \in C$ (with parameters W_i) performs, based on the local validation sets.
- Step 3b: A new function $f(w_1, w_2, \dots, w_N)$ is considered to compute a parameter $w_G \in W_G$ of the global model from the corresponding parameters $w_i \in W_i$ in the local models of all clients $i \in C$. This function is based on how well the model of each client performs in Step 3a.

Data heterogeneity is commonly observed in classic FL setups. This means that, often, clients with larger datasets have a bigger footprint on the global model. However, instead of associating each model with an importance proportional to the size of their datasets, the proposed aggregation procedure computes a local model’s contribution to the global model based on its performance on the local validation dataset(s).

Step 3a: Evaluating the performance of the local models

In Step 2, all clients send their local model to the federated learning server after training on their local datasets. In Step 3a, the server is now interested in evaluating each local model’s performance on

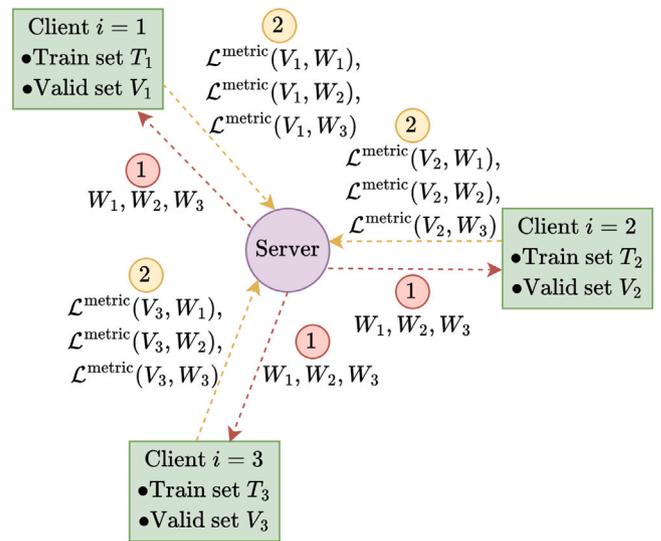


Fig. 3. Schematic illustration of the full validation policy, where (1) denotes the first step and (2) the second step in the communication sequence.

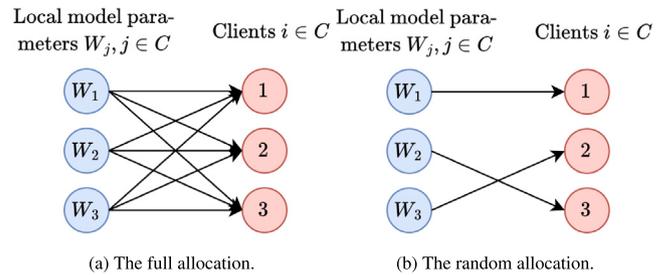


Fig. 4. An example of a full and random allocation with three clients.

another client’s local validation dataset. As such, the server coordinates the clients that shall evaluate each local model, and distributes these models to these clients. Since the server is only interested in the model’s performance on another client’s dataset, each client provides the loss their assigned model achieved against its validation dataset. The fact that the server is coordinating the distribution of the local models also limits the possibility of identifying the client that created the model under evaluation.

(i) *full validation policy.* Under the full validation policy, the server sends for each client $j \in C$, the parameters W_j of the local model of client j to all clients $i \in C$ (see Fig. 4(a)). Each client $i \in C$ then computes the loss $\mathcal{L}^{\text{metric}}(V_i, W_j)$ of the local model of client j , with parameters W_j , on its own validation set V_i using a standard metric. This metric could be, for instance, the (Root) Mean Squared Error (for regression problems), or the Cross Entropy loss (for classification problems). Each client $i \in C$ then sends this loss $\mathcal{L}^{\text{metric}}(V_i, W_j)$ of the local model of client j back to the server. This full validation policy is illustrated in Fig. 3.

The server then computes the evaluation score E_j of client j , that represents the accuracy of the local model of client j . This evaluation score is the median of all losses $\mathcal{L}^{\text{metric}}(V_i, W_j)$ over all clients $i \in C$:

$$E_j = \text{median}_{i \in C} \{ \mathcal{L}^{\text{metric}}(V_i, W_j) \}. \quad (4)$$

Taking the median of all losses, rather than the average or sum, decreases the influence of clients with noisy data on the evaluation score. A low evaluation score E_j for a client j implies that the model

of client j performs well. The pseudocode for the full validation policy is in Algorithm 1.

Algorithm 1: Step 3a - The full validation policy

```

1 for Client  $j \in C$  do
2   for Client  $i \in C$  do
3     Send the parameters  $W_j$  to client  $i$ ;
4     Calculate the loss  $\mathcal{L}^{\text{metric}}(V_i, W_j)$  with i) the parameters
        $W_j$  of the local model of client  $j$  and ii) the validation
       set  $V_i$  of client  $i$ ;
5     Send the loss  $\mathcal{L}^{\text{metric}}(V_i, W_j)$  back to the server;
6   end
7   Calculate the evaluation score:
        $E_j = \text{median}_{i \in C} \{ \mathcal{L}^{\text{metric}}(V_i, W_j) \}$ ;
8 end

```

(ii) *random validation policy*. The full validation policy is relatively time-consuming, since the validation loss of all local models is computed with the local validation set of all clients. The random validation policy therefore computes the loss of each local model on only one validation set V_i at a client $i \in C$. This random validation policy is illustrated in Fig. 5.

The pseudocode for the random validation policy is in Algorithm 2. The server first randomly allocates one local model of a client j , with parameters W_j , to each client $i \in C$. Each local model (with parameters W_j) is thus allocated to exactly one client $i \in C$, and reversely each client $i \in C$ has one local model (with parameters W_j) allocated to it. Let $i_j^{\text{ass}} \in C$ denote the client to which the local model of client j is assigned. An illustration of such an allocation with three clients is depicted in Fig. 4(b).

For each client $j \in C$ the loss $\mathcal{L}^{\text{metric}}(V_{i_j^{\text{ass}}}, W_j)$ of the local model (with parameters W_j) is computed on the validation set $V_{i_j^{\text{ass}}}$ of the assigned client $i_j^{\text{ass}} \in C$. The evaluation score of client j is then simply this loss:

$$E_j = \mathcal{L}^{\text{metric}}(V_{i_j^{\text{ass}}}, W_j). \quad (5)$$

Algorithm 2: Step 3a - The random validation policy

```

1 Randomly assign the local model of each client  $j \in C$  to one
   client  $i \in C$ . Ensure that exactly one local model is assigned to
   each client  $i \in C$ . Let  $i_j^{\text{ass}} \in C$  denote the client to which the
   local model of client  $j$  is assigned;
2 for Client  $j \in C$  do
3   Send the parameters  $W_j$  to the assigned client  $i_j^{\text{ass}}$ ;
4   Compute the loss  $\mathcal{L}^{\text{metric}}(V_{i_j^{\text{ass}}}, W_j)$  with i) the parameters
        $W_j$  of the local model of client  $j$  and ii) the validation set
        $V_{i_j^{\text{ass}}}$  of client  $i_j^{\text{ass}}$ ;
5   Send the loss  $\mathcal{L}^{\text{metric}}(V_{i_j^{\text{ass}}}, W_j)$  back to the server;
6   Set the evaluation score:  $E_j = \mathcal{L}^{\text{metric}}(V_{i_j^{\text{ass}}}, W_j)$ ;
7 end

```

Step 3b: Aggregating the model parameters

In step 3b, the parameters of the global model W_G are computed, based on the parameters W_i and the evaluation scores E_i of the local models of the clients $i \in C$. Here, two methods to aggregate the parameters are proposed: (i) The best model aggregation policy, which selects the best model of all local models as the new global model and (ii) the softmax aggregation policy, which uses a weighted average to compute the parameters of the global model.

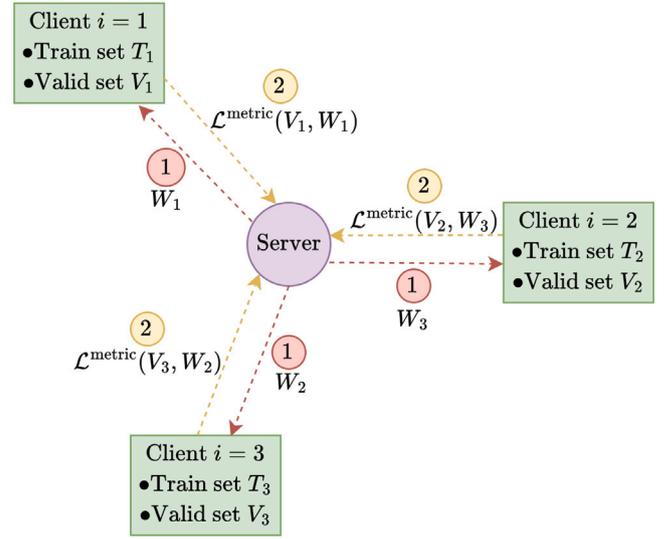


Fig. 5. Schematic illustration of the random validation policy. In this example, the weights of the clients are allocated as in Fig. 4(b). (1) denotes the first step and (2) the second step in the communication sequence.

Best model aggregation policy. The best model aggregation policy is an adaptation of the work of Blanchard et al. [14] that selects the best performing model based on an evaluation heuristic. The adaptation considered in this paper consists of modifying the evaluation procedure to use the random and full validation policies presented in Step 3a. As a result, based on the evaluation score provided by the validation policy, the model that has the best evaluation score is selected as the new global model, as follows:

$$i^{\text{best}} = \arg \min_{i \in C} \{ E_i \}, \quad (6)$$

$$W_G = W_{i^{\text{best}}}, \quad (7)$$

where i^{best} is the client with the lowest evaluation score.

After each epoch, the weights of only a single model are selected as weights of the global model. However, the client with the best model differs from epoch to epoch. All clients, that are the client with the best model at least once, thus contribute to the final global weights through the weight updating algorithm. This is illustrated in Fig. 6 for the standard gradient descent algorithm. Here, γ denotes the learning rate, and m_i^e denotes the gradient of epoch e at client i . After epoch 1, client 1 is selected as client with the best model, and the global weights are updated with the gradients of client 1. In epoch 2, however, client 3 is selected as client with the best model, and the gradients of client 3 are used to further update the weights. At epoch 3, the gradients of client 2 are used to update the weights. In this way, all three clients contribute to the final weights of the model.

Softmax aggregation policy. For the softmax policy, the parameters of the global model are computed as a weighted average of the parameters of the local model. In this average, each client $i \in C$ gets a weight α_i , based on the evaluation score E_i .

A local model performs better if the evaluation score is lower. First, the inverse of the evaluation score A_i for each client $i \in C$ is therefore computed:

$$A_i = \frac{1}{E_i} \quad (8)$$

Then, the evaluation scores are normalized to diversify the final weights, after the softmax function. First, the mean μ and the standard deviation σ of the inverse evaluation scores are computed:

$$\mu = \frac{1}{N} \sum_{i \in C} A_i \quad (9)$$

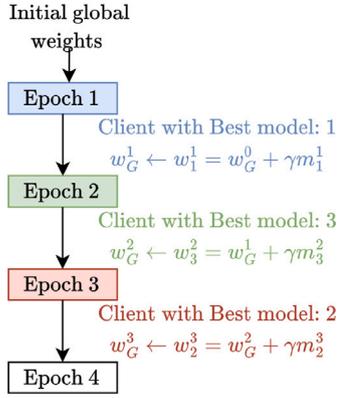


Fig. 6. Example of the *Best model* aggregation policy with the gradient descent algorithm. Here, γ denotes the learning rate, w_i^e denotes a single model parameter at epoch e and client i , w_G^e denotes the corresponding single parameter weight at epoch e for the global model, and m_i^e denotes the gradient of this parameter at epoch e and client i .

$$\sigma = \sqrt{\frac{1}{N-1} \sum_{i \in C} (A_i - \mu)^2} \quad (10)$$

With this, the z-score Z_i for each client $i \in C$ is computed:

$$Z_i = \frac{A_i - \mu}{\sigma} \quad (11)$$

Last, the weight α_i of each client $i \in C$ with the softmax function is computed:

$$\alpha_i = \frac{\exp(Z_i)}{\sum_{j \in C} \exp(Z_j)}. \quad (12)$$

The softmax function ensures that each weight is between 0 and 1, and that the sum of the weights is 1 (i.e., $\sum_{i \in C} \alpha_i = 1$). Ultimately, a parameter $w_G \in W_G$ of the global model is computed with by taking the weighted average of the corresponding parameter $w_i \in W_i$ of the local models of the clients $i \in C$:

$$w_G = \sum_{i \in C} \alpha_i w_i. \quad (13)$$

The Z-score is a common normalization technique [42], and is used in this paper to obtain weights that better reflect the difference in evaluation scores (E_i). As an example, if airline 1's model has the validation score $E_1 = 20$, and thus $A_1 = \frac{1}{20} = 0.05$, and airline 2's model has the validation score $E_2 = 25$ and thus $A_2 = \frac{1}{25} = 0.04$, then the weight attributed to each model using A_1 and A_2 directly as input to the softmax function, is respectively 0.5025 and 0.4975. In this case both models have an approximately equal contribution to the global model, despite the significant difference in their evaluation scores. To avoid this, it is important to normalize the range obtained by the evaluation score using Eq. (11) corresponding to the Zscore computation. As a result, after computing the mean and standard deviation, $Z_1 = 0.707$ and $Z_2 = -0.707$, and using these values as input to the softmax function, the airline models resulting weights are $\alpha_1 = 0.804$ and $\alpha_2 = 0.196$. As such, using the Z-score normalization has allowed a better differentiation of the weights attributed to each airline's models, which is proportional to their relative performance.

Novel aggregation methods for the global model parameters

The two validation policies (Full validation and Random validation) and two aggregation policies (Best Model and Softmax) are proposed. The combination of these policies give four novel aggregation methods, as specified in Table 1. These aggregation methods all replace step 3 of the FL framework in Section 3.

Table 1

The four proposed aggregation methods based on the validation policy and the aggregation policy.

Name method	Validation policy	Aggregation policy
Random-Best	Random	Best model
Random-Softmax	Random	Softmax
Full-Best	Full	Best model
Full-Softmax	Full	Softmax

5. Case study: A multi-airline federated learning framework for RUL prognostics

The proposed FL framework with a decentralized validation procedure is illustrated for the case of airlines collaborating to generate a common machine learning model for RUL prognostics of aircraft engines. An airline operates several aircraft. The engines of these aircraft are continuously monitored due to their safety criticality. The obtained measurements are used to learn degradation patterns and anticipate engine failures. To enhance this learning, airlines are interested in learning from other airlines' datasets, provided that the data privacy is preserved.

The proposed FL framework is illustrated using the condition monitoring data samples of aircraft engines in folder D502 of the N-CMAPSS dataset [23]. The N-CMAPSS dataset consists of both sensor measurements and operating conditions of several engines. These measurements are available from the first flight performed with the engine, until the failure of the engine. Dataset D502 consists of six engines (engines 2, 5, 10, 16, 18 and 20) whose data is meant for training, and three engines (engines 11, 14 and 15) to be used for testing.

In this study, a case is considered where each airline has the data of one engine. Specifically, the training engines are available at Airlines A-F as follows: Airline A (engine 2), Airline B (engine 5), Airline C (engine 10), Airline D (engine 16), Airline E (engine 18), Airline F (engine 20). The data of these six training engines is used to train a RUL prognostic model. The test engines 11, 14, 15 are not assigned to a specific airline. They are instead used to analyze the performance of the RUL prognostic model by predicting the RUL of the test engines 11, 14 and 15 after each flight. With this, the performance of the proposed FL framework and the robust aggregation methods are analyzed.

Table 3 gives a detailed overview of the considered airlines and their engines. The flights of each engine belong to one flight class: flight class 1 (short flights), flight class 2 (medium long flights) or flight class 3 (long flights).

5.1. Data preprocessing

The N-CMAPSS dataset contains the measurements of 28 different sensors. However, many of the measurements of these sensors are highly correlated. Following [3], a total of 13 sensors (see Table 2) are selected as input for the RUL prognostic model based on the correlation. After this sensor selection, 89 million sensor measurements are still available. Previous studies using N-CMAPSS [3,43] aggregate the measurements to reduce the computational time to train a neural network. Similarly, as measurement aggregation, the mean measurement per sensor and per operating condition per 20 s is considered in this study. This reduces the dataset to over 4 million measurements.

Lastly, the measurements are normalized using min-max normalization. The measurements of a training engine (i.e., 2, 5, 10, 16, 18, 20) are normalized using the minimum and maximum measurement belonging to only this engine. In this way, airlines do not share their minimum and maximum measurement with each other. The measurements of the test engines (i.e., 11, 14, 15), however, are normalized with the minimum and maximum measurement from all training engines (i.e., 2, 5, 10, 16, 18, 20), since these engines are not assigned to a specific airline.

Table 2

Selected sensors and operating conditions which are used as input to the RUL prognostic model, LPC - Low Pressure Combustor. HPC - High Pressure Combustor. HPT - High Pressure Turbine. LPT - Low Pressure Turbine.

Symbol	Description	Unit
Selected sensors		
Wf	Fuel flow	pps
Nf	Physical fan speed	rpm
T24	Total temperature at LPC outlet	°R
T30	Total temperature at HPC outlet	°R
T48	Total temperature at HPT outlet	°R
T50	Total temperature at LPT outlet	°R
P2	Total pressure at fan inlet	psia
P50	Total pressure at LPT outlet	psia
W21	Fan flow	pps
W50	Flow out of LPT	lbm/s
SmFan	Fan stall margin	-
SmLPC	LPC stall margin	-
SmHPC	HPC stall margin	-
Operating conditions		
Mach	Flight Mach number	-
TRA	Throttle-resolver angle	%
alt	Altitude of the aircraft	ft
T2	Total temperature at the fan inlet	°R

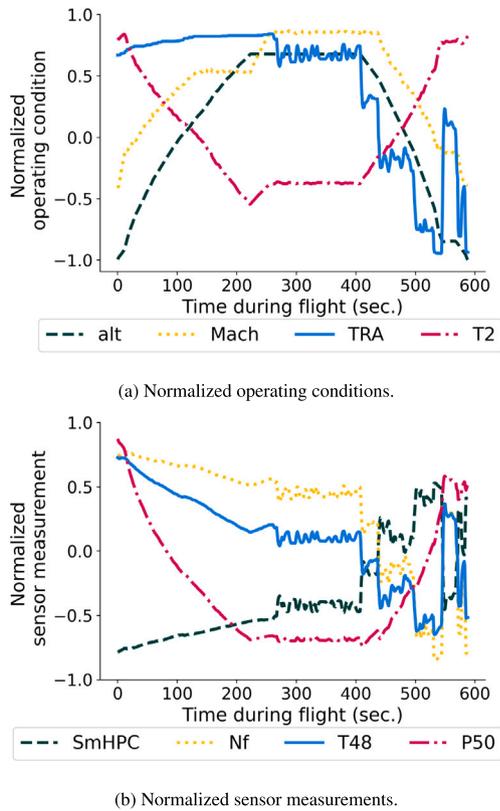


Fig. 7. Normalized operating conditions and normalized sensor measurements of flight 1, training engine 2.

After normalization, a validation/training data split is performed for each training engine considered in the dataset: 20% of the flights of each training engine are randomly selected as validation data, and the remaining 80% of the flights of this engine become training data.

Fig. 7 shows an example of the considered measurements. **Fig. 7(a)** shows the aggregated, normalized measurements of the four operating conditions during the first flight of training engine 2. **Fig. 7(b)** shows the aggregated normalized measurements of four sensors during the same flight. Due to the min-max normalization, all measurements are

between -1 and 1 . The operating conditions and sensor measurements clearly show the flight profile. For instance, the altitude of the aircraft (“alt”) first increases during the take-off, is then stable during cruise, and decreases when the aircraft is landing. Similarly, the pressure at the Low Pressure Turbine outlet (“P50”) is high during the take-off and landing, and low during cruise.

Let $\mathbf{X}^{e,f} = [\mathbf{X}_1^{e,f}, \mathbf{X}_2^{e,f}, \dots, \mathbf{X}_{M^{e,f}}^{e,f}]$ denote the normalized multi-sensor measurements during flight f for an engine e . Here, $M^{e,f}$ is the number of measurements from this flight. Each normalized multi-sensor measurement $\mathbf{X}_g^{e,f}$ (with $g \in \{1, 2, \dots, M^{e,f}\}$) contains H measurements, i.e., $\mathbf{X}_g^{e,f} = [x_g^{e,f,1}, x_g^{e,f,2}, \dots, x_g^{e,f,H}]$. Also, $x_g^{e,f,h}$ (with $h \in \{1, 2, \dots, H\}$) is the g th normalized measurement of sensor h during flight f of engine e . Since 13 sensors and 4 operating conditions are considered in this case study, there are $H = 17$ different measurements in total.

The number of measurements $M^{e,f}$ varies per engine e and per flight f . The RUL prognostic model, however, can only handle input samples of a fixed dimension. To address this, for each engine e and each flight f , several samples of length 50 are extracted from $\mathbf{X}^{e,f}$. The first sample $\tilde{\mathbf{X}}_1^{e,f}$ of flight f of engine e consists of the first 50 multi-sensor measurements of $\mathbf{X}^{e,f}$, i.e., $\tilde{\mathbf{X}}_1^{e,f} = [\mathbf{X}_1^{e,f}, \mathbf{X}_2^{e,f}, \dots, \mathbf{X}_{50}^{e,f}]$. Moving forward with a step size of 10 time-steps, the second sample $\tilde{\mathbf{X}}_{11}^{e,f}$ of length 50 from this flight is created as $\tilde{\mathbf{X}}_{11}^{e,f} = [\mathbf{X}_{11}^{e,f}, \mathbf{X}_{12}^{e,f}, \dots, \mathbf{X}_{60}^{e,f}]$. Again moving 10 time-steps, the third sample $\tilde{\mathbf{X}}_{21}^{e,f}$ of this flight is created. This is repeated until the end of the flight. The data sample of engine e , flight f and time-step s is denoted by $\tilde{\mathbf{X}}_s^{e,f} = [\mathbf{X}_s^{e,f}, \mathbf{X}_{s+1}^{e,f}, \dots, \mathbf{X}_{s+50}^{e,f}]$. Let $S^{e,f}$ be the set with all the time-steps s belonging to flight f of engine e .

5.2. Convolutional neural network for RUL prognostics

The RUL of the aircraft engines is predicted using an one-dimensional Convolutional Neural Network (1D-CNN). CNNs are originally developed for image data. However, the ability of CNNs to extract features from the data makes them also very suitable for RUL prediction from time-series of sensor measurements [44]. Especially 1D-CNNs have been frequently used for RUL prediction using multi-variate time series, where their ability to extract features yielded state-of-the-art results [44–46].

Fig. 8 shows a schematic overview of the considered CNN. The CNN consists of 3 convolutional layers. The first 2 convolutional layers consist of 10 kernels, each of size 9×1 (i.e., one-dimensional kernels), that move with a stride of 1. The last convolutional layer has one kernel, and compresses all 10 feature maps into one single feature map. The convolutional operation in the c th convolutional layer for the p th kernel k_p^c is [45]:

$$z_p^c = \text{ReLU}(k_p^c * z^{c-1} + b_p^c) \quad (14)$$

where z_p^c is the p th feature map of layer c , with b_p^c the corresponding bias of this feature map, z^{c-1} represents all feature maps in layer $c-1$, $*$ is the convolutional operator and $\text{ReLU}(\cdot)$ is the considered activation function, namely the Rectified Linear Unit.

After the convolutional layers, the CNN consists of two fully connected layers. The first fully connected layer consists of 100 nodes. Let z^C be the feature map from the last convolutional layer C . The output z^l of the first fully connected layer is then [45]:

$$z^l = \text{ReLU}(w^l z^C + b^l), \quad (15)$$

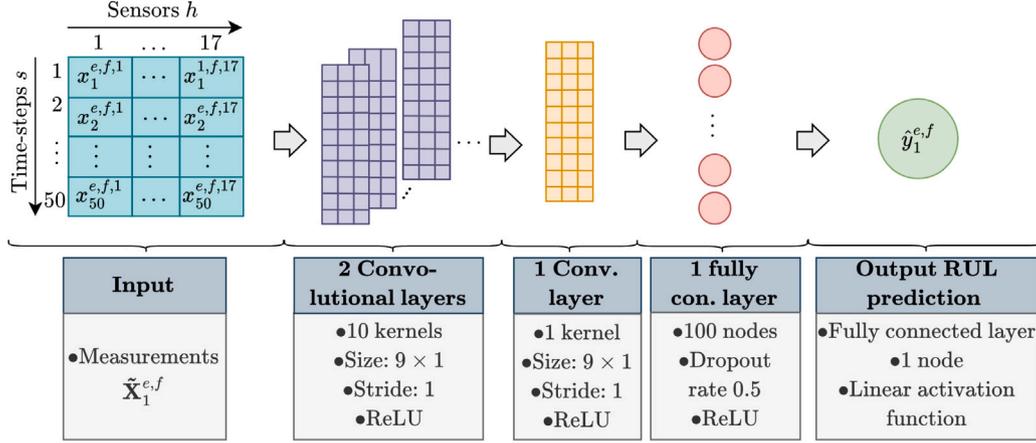
with b^l the bias and w^l the weights of the fully connected layer. A dropout rate of 0.5 is applied on the output of this layer to prevent overfitting. Lastly, the final layer contains one node, and outputs the RUL prognostic using a linear activation function. The CNN is trained with the Adam optimizer [47] for 100 epochs with a batch size of 128 and a learning rate of 0.001.

The Root Mean Squared Error (RMSE) between the true RUL and the estimated RUL is used as the loss function to train this CNN. With

Table 3

An overview of the airlines and their engines. Airlines A-F each have one engine for which measurements are recorded until the actual failure of each engine. Fault mode 1: Degradation in the efficiency of the High Pressure Turbine (HPT). Fault mode 2: Degradation in the efficiency of the HPT, and in the efficiency and flow of the Low Pressure Turbine (LPT).

Airline (engine)	Training						Testing		
	A (2)	B (5)	C (10)	D (16)	E (18)	F (20)	x (11)	x (14)	x (15)
Flight class	3	3	3	3	3	3	3	1	2
Fault mode	1	1	1	2	2	2	2	2	2
Number of flights	75	89	82	63	71	66	59	76	67
Mean length of flight (s)	11,375	11,611	11,618	12,148	12,545	11,639	11,246	2063	6470

**Fig. 8.** Schematic overview of the CNN with a data sample $\tilde{X}_1^{e,f}$.

the new aggregation methods, the RMSE is also used as the metric loss $\mathcal{L}^{\text{metric}}(V_i, W_j)$ to compute the evaluation score E_j of a local model of an airline $j \in C$ in \mathcal{V}_i . To validate the global model, with the parameters W_G (Step 5 in the FL framework, Section 3), the sum of the loss is computed for each airline $i \in C$ as the sum of the squared errors:

$$\mathcal{L}^{\text{sum}}(V_i, W_G) = \sum_{j=1}^{|V_i|} (\hat{y}_j - y_j)^2, \quad (16)$$

where y_j is the true RUL of the j th sample in the validation set V_i , and \hat{y}_j is the estimated RUL of this sample.

As described in Section 5.1, the data of each flight f of engine e is divided into multiple data samples, where each data sample $\tilde{X}_s^{e,f}$ belongs to a time-step s of flight f of engine e . The set $S^{e,f}$ contains all time-steps s belonging to flight f of engine e . Each data sample $\tilde{X}_s^{e,f}, s \in S^{e,f}$ is used as input to the CNN, to obtain as output the predicted RUL $\hat{y}_s^{e,f}$. The final estimated RUL $\hat{y}^{e,f}$ of engine e after flight f is the median of the RUL prognostic of all samples belonging to that flight:

$$\hat{y}^{e,f} = \text{Median}_{s \in S^{e,f}} (\hat{y}_s^{e,f}). \quad (17)$$

5.3. Benchmark models

The FL framework is compared with two benchmark models: the Unrestricted access Centralized (UC) learning model and the Non-collaborative Isolated (NI) learning model.

Unrestricted access Centralized (UC) learning model

For this model, it is envisioned that all airlines fully cooperate by sharing all their data, without any privacy barriers. A central server trains a RUL prognostic model on engines 2, 5, 10, 16, 18 and 20 and tests this model on engines 11, 14 and 15. This is an idealistic case when all data is shared. The outcome of this benchmark is a lower bound on the performance of the FL framework.

For the UC learning model, the CNN is trained with all the available data of all six training engines. The sensor measurements and operating

conditions are normalized using min-max normalization. Since all data is available at a central server, the minima and maxima across all data from all training engines is used for normalization. The training/validation split after normalization remains the same as in the FL framework.

Non-collaborative isolated (NI) learning model

For the NI model, it is envisioned that there is no collaboration between airlines. Each airline trains its own RUL prognostic model with the data of the one engine it has available. The considered cases are as follows: Airline A (engine 2) trains a CNN with the data of engine 2, Airline B (engine 5) trains a CNN with the data of engine 5, Airline C (engine 10) trains a CNN with the data of engine 10, Airline D (engine 16) trains a CNN with the data of engine 16, Airline E (engine 18) trains a CNN with the data of engine 18, and Airline F (engine 20) trains a CNN with the data of engine 20 (see Table 3). In total, six different CNNs are trained, one CNN per airline. As in the FL framework, the measurements of a training engine are normalized using the minimum and maximum measurement belonging to only this engine. After normalization, the training/validation split remains the same as in the FL framework.

The performance of the proposed FL framework vs. the NI learning model is analyzed by testing these six CNNs on the three test engines 11, 14 and 15. In contrast to the FL framework, when testing the CNN of airline $i, i \in \{A, B, C, D, E, F\}$, the measurements of the test engines are normalized with the minimum and maximum measurement from the one training engine belonging to airline $i, i \in \{A, B, C, D, E, F\}$.

5.4. Simulation of noisy condition-monitoring data

In Section 4, four novel aggregation methods are proposed to develop a global model that is resilient to clients with noisy data. To evaluate these aggregation methods in the presence of noisy data, a scenario is simulated where some airlines have noisy data. Here, a varying amount of noise is added to both the sensor measurements and the operating conditions of the engine assigned to these airlines.

Table 4

Comparison of the RMSE and MAE (both in flights) of the FL model and the Unrestricted Access centralized (UC) learning model. The lowest obtained RMSE for the validation set, and the number of epochs after which this lowest RMSE is obtained, are also given. The best results are denoted in bold.

Model	Validation results		Test results							
	Min. valid.	Obt. after	Engine 11		Engine 14		Engine 15		All engines	
	RMSE	... epochs	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
FL model	12.7	61	6.2	5.0	13.5	9.2	7.4	5.3	9.9	6.7
UC model	9.8	85	6.3	5.0	8.0	5.6	3.7	2.4	6.3	4.4

Table 5

Comparison of the RMSE and MAE (both in flights) of the FL model and the six non-collaborative isolated (NI) learning models (each trained with only the training data of one engine). The best results are denoted in bold.

	Test results							
	Engine 11		Engine 14		Engine 15		All engines	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
FL model								
	6.2	5.0	13.5	9.2	7.4	5.3	9.9	6.7
Six NI learning models								
Airline A (Engine 2)	22.9	20.0	22.3	19.1	21.1	19.2	22.1	19.4
Airline B (Engine 5)	8.4	6.5	16.3	11.6	10.1	7.0	12.4	8.6
Airline C (Engine 10)	9.6	7.4	25.7	21.1	17.2	13.7	19.3	14.6
Airline D (Engine 16)	6.8	5.6	10.8	8.0	7.0	5.6	8.6	6.9
Airline E (Engine 18)	18.8	15.9	17.2	14.3	15.9	13.7	17.3	14.6
Airline F (Engine 20)	10.6	8.9	16.9	13.3	16.4	13.8	15.2	12.2
Mean over all NI learning models	12.9	10.7	18.2	14.7	14.6	12.2	15.8	12.7

Assume that engine e_i is available at airline $i \in C$, and that this airline i is selected such that noise is added to the measurements. Noise is added *before* normalizing the data. For each non-normalized measurement $\tilde{x}_g^{e_i, f, h}$ (i.e., the g th non-normalized measurement of sensor h during flight f of engine e_i), noise is sampled from a normal distribution with a mean of 0 and a standard deviation of $\alpha \cdot \sigma_{i, h}$:

$$\tilde{x}_g^{e_i, f, h} \leftarrow \tilde{x}_g^{e_i, f, h} + \mathcal{N}(0, \alpha \cdot \sigma_{i, h}), \quad (18)$$

$$\forall f \in \{1, \dots, F_{e_i}\}, \forall g \in \{1, \dots, M^{e_i, f}\}, \forall h \in \{1, \dots, H\}.$$

Here, α is a user-chosen parameter, F_{e_i} is the number of flights of engine e_i and $\sigma_{i, h}$ is the standard deviation of all non-normalized measurements of a sensor h for all flights of the considered engine e_i .

6. Results

6.1. Results for the FL (Federated Learning) model vs. the UC learning (centralized) model

In this section, the results achieved with the proposed FL model are presented. These results are compared with the unrestricted access centralized (UC) learning model, where the ideal case in which all airlines share their data to train a common RUL prognostic model is considered (see Section 5.3).

Table 4 shows the minimum RMSE of the validation set for both models. Fig. 9 shows the RMSE of the validation set after each epoch. The RMSE of the validation set is consistently lower for the UC learning model, than for the FL model. This is expected, since the centralized model is directly trained with all the data from all the airlines. For the FL model, the minimum validation loss is already obtained after 61 epochs, while the model is trained for 100 epochs. At the end of the training procedure, the model might thus overfit. However, with the decentralized validation procedure, the weights of the model after 61 epochs are identified as the best weights to generate the test results with.

Table 4 also shows the RMSE and Mean Absolute Error (MAE) of the RUL prognostics for the three test engines, while Fig. 10 shows the corresponding RUL prognostic over time. Also here, the UC learning model outperforms the FL model, with an equal or lower RMSE and

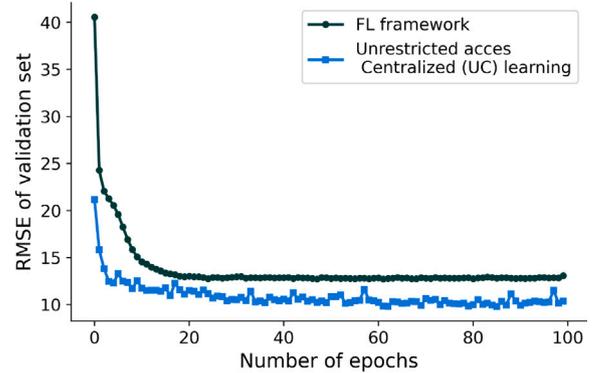


Fig. 9. RMSE of the validation set for the FL model and the unrestricted access centralized (UC) learning model.

MAE of the RUL prognostics for each test engine, except the RMSE of test engine 11. However, the overall RMSE of the FL model (9.9 flights) is close to the overall RMSE of the UC learning model (6.3 flights).

6.2. Results for the FL model vs. the NI learning (isolated) models

In this section, the results achieved with the proposed FL model are compared with the results achieved with the non-collaborative (NI) learning models. The aim of this comparison is to determine whether a better RUL prognostic model is obtained if airlines collaborate through the FL framework instead of each training their own prognostic model (NI learning, see Section 5.3).

Table 5 shows the RMSE and MAE of the RUL prognostics for the test engines. Here, the FL model has a lower overall RMSE (9.9 flights) and MAE (6.6 flights) than the mean RMSE (15.8 flights) and mean MAE (12.7 flights) over all six NI learning models. Moreover, the FL model consistently outperforms five out of the six NI learning models, with a lower RMSE and MAE for all three test engines.

However, the overall RMSE of the model of airline D, trained with engine 16, is 8.6 flights, which is lower than the overall RMSE of

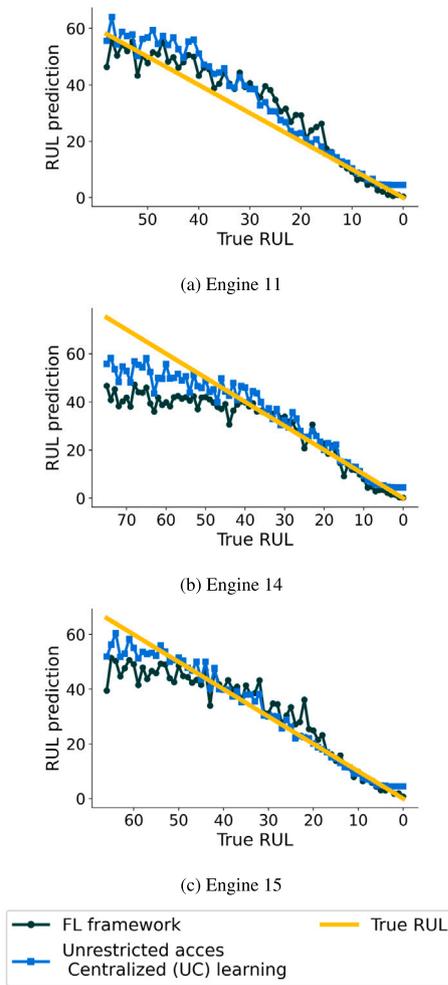


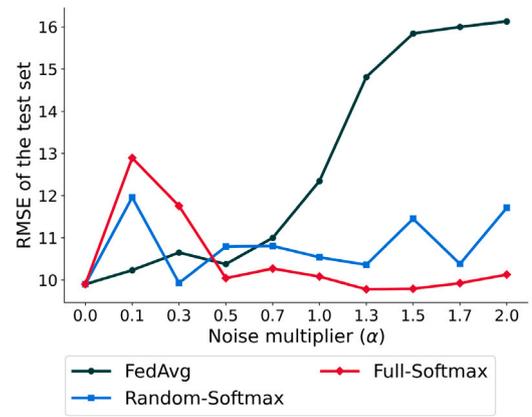
Fig. 10. RUL prognostics per engine for the FL model and the Unrestricted Access centralized (UC) learning model.

the FL model (9.9 flights). This illustrates the heterogeneity in the experimental setup, where Airline D possesses an engine that has similar failure conditions and fault modes as those observed in the test dataset (see Table 5), while Airlines A, B, C, E and F do not benefit from such a representative dataset. However, an airline does not know beforehand which types of failures will occur in the future. As the FL model provides the best results for five out of six airlines, an airline is likely to obtain better results when participating in the FL model. In conclusion, it is therefore indeed beneficial for an airline to train a common RUL prognostic model through the proposed FL framework.

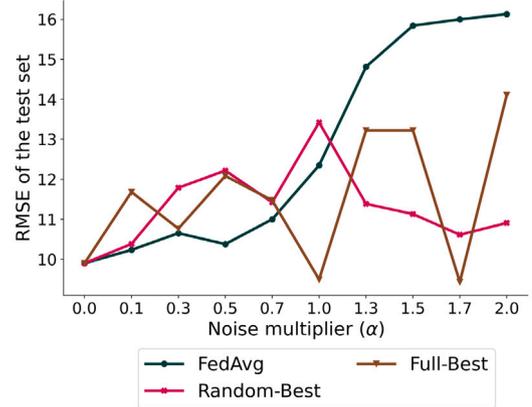
6.3. Comparison of the novel aggregation methods for the global model parameters

In this section, noise is added to the sensor measurements and operating conditions of training engine 5 and 18 (see Section 5.4). Here, training engine 5 fails according to fault mode 1, while training engine 18 fails according to fault mode 2 (Section 5). The size of the noise multiplier α is varied from 0 (no additional noise) to 2 (large additional noise). Next, the four proposed aggregation methods of Section 4 are compared with the standard FedAvg [7] aggregation method (Section 3).

Fig. 12 shows a plot of both the normalized and the non-normalized measurements of four selected sensors during flight 1, engine 2, with added noise with a multiplier of $\alpha = 0$, $\alpha = 0.1$ and of $\alpha = 1$.



(a) The RMSE of the test set with FedAvg, Random-Softmax and Full-Softmax.



(b) The RMSE of the test set with FedAvg, Random-Best and Full-Best.

Fig. 11. The overall RMSE of the test set in flights for all considered aggregation methods, while considering different noise multipliers α . Here, additional noise is added to the data of engine 5 and 18.

Table 6

Comparison of the RMSE and MAE (in flights) of the RUL prognostics for all three test engines in N-CMAPSS, for all considered aggregation methods. Here, no noise is added to the data (i.e., $\alpha = 0$). The best results are denoted in bold.

Aggregation Method	Engine 11		Engine 14		Engine 15		All engines	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
FedAvg	6.2	5.0	13.5	9.2	7.4	5.3	9.9	6.7
Softmax aggregation policy								
Random-Softmax	8.6	7.0	15.1	10.2	7.7	5.6	11.2	7.7
Full-Softmax	8.3	6.9	12.6	9.3	8.1	6.6	10.1	7.7
Best model aggregation policy								
Random-Best	11.2	9.4	13.6	12.0	11.3	9.9	12.2	10.6
Full-Best	8.1	6.5	15.2	10.6	9.2	6.8	11.6	8.1

Table 6 shows the RMSE and MAE of the RUL prognostics of the three test engines, for all considered aggregation methods, when there is no noise added to the data. The FedAvg method provides the best results, but the results with the softmax aggregation policy are only slightly worse. Full-Softmax and Random-Softmax have the same MAE for all test engines, but the RMSE of all test engines is considerably lower with Full-Softmax. The best model aggregation policy (Random-Best and Full-Best) gives a relatively high RMSE and MAE.

Fig. 11 shows the RMSE of the test engines for a noise multiplier α ranging from 0 (no noise) to 2 (very large noise). Tables 7 and 8

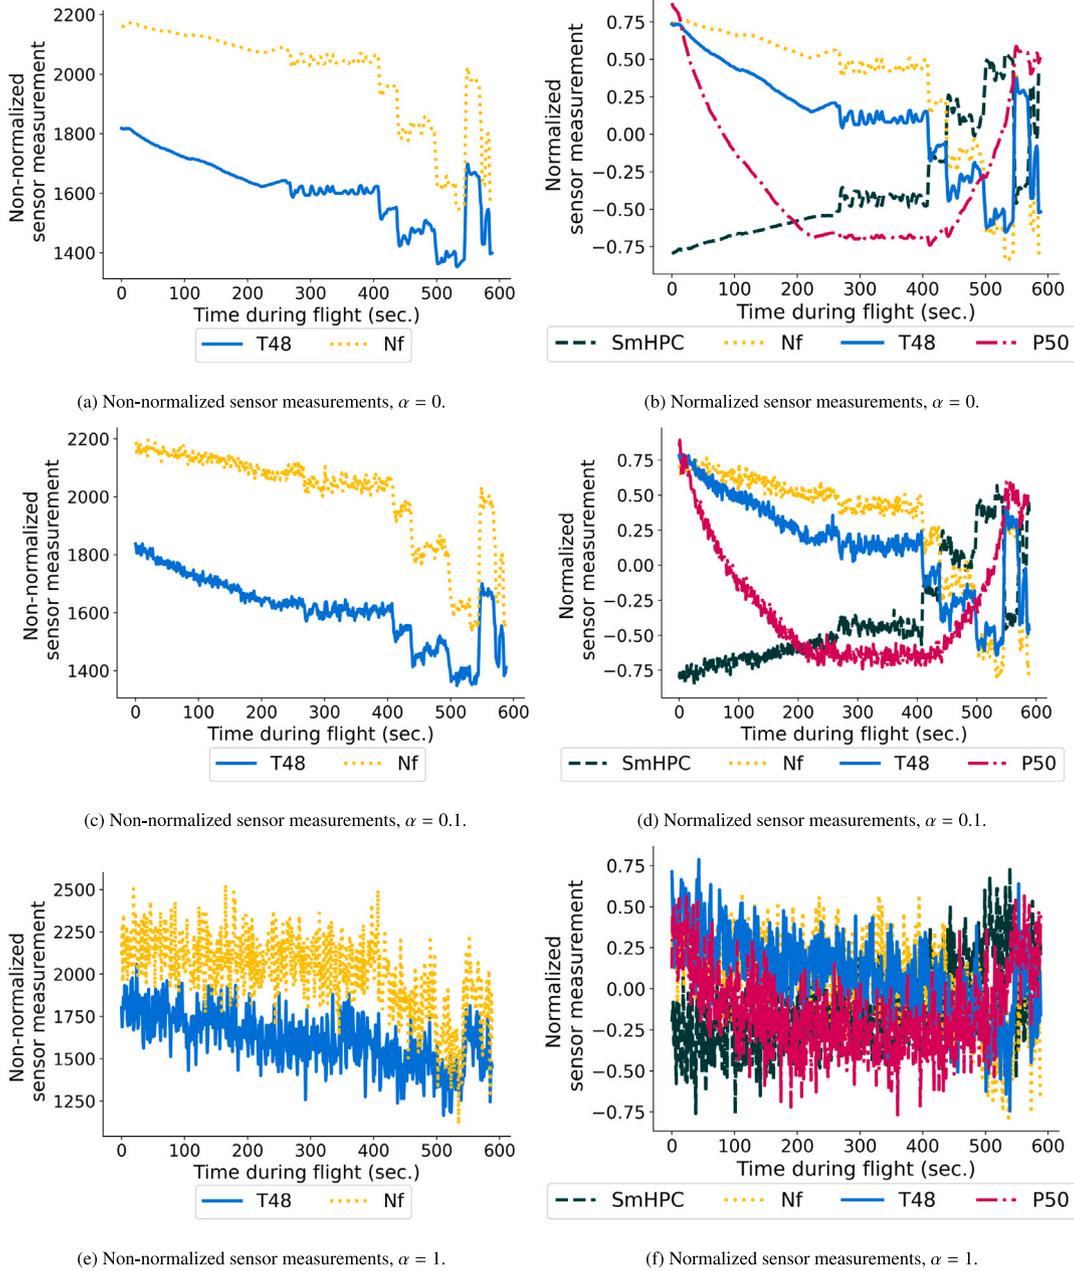


Fig. 12. Non-normalized and normalized sensor measurements of flight 1, engine 2, for different values of the noise multiplier α . For the non-normalized sensor measurements, only sensor T48 and Nf are plotted, since the scales of the other sensors are too different.

Table 7

Comparison of the RMSE and MAE (in flights) of the RUL prognostics for all three test engines in N-CMAPSS, for all considered aggregation methods. Here, additional noise with a standard deviation multiplier of $\alpha = 0.1$ is added to the data of engine 5 and 18. The best results are denoted in bold.

Aggregation Method	Engine 11		Engine 14		Engine 15		All engines	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
FedAvg	8.8	7.2	12.8	9.2	7.8	6.2	10.2	7.6
Softmax aggregation policy								
Random-Softmax	6.0	4.8	17.2	12.5	8.0	5.8	12.0	8.0
Full-Softmax	5.2	4.0	18.9	14.2	8.6	6.0	12.9	8.5
Best model aggregation policy								
Random-Best	6.8	5.8	13.5	10.4	8.7	7.1	10.4	7.9
Full-Best	5.6	4.5	16.2	11.2	9.3	6.5	11.7	7.7

show the corresponding test results for a noise multiplier of $\alpha = 0.1$ and $\alpha = 1$ respectively. The FedAvg method performs well when the noise is relatively small, i.e., when $\alpha \leq 0.5$. For these noise values, the RMSE of all test engines is quite small. For $\alpha = 0$ (no noise) and $\alpha = 0.1$ (small noise), the RMSE is even smaller with the FedAvg method, than with all proposed aggregation methods. When the noise increases, however, i.e., from $\alpha = 0.7$ onwards, the RMSE of the test engines increases as well (see Fig. 11).

With the Softmax aggregation policies, the RMSE over all test engines is relatively high when α is small (see Fig. 11(a)). Table 7 shows that this is mainly because for small α and the Softmax aggregation policies, the predictions for test engine 14 are inaccurate. When α increases, however, the RMSE over all test engines remains stable, and even decreases compared to the smaller values of α . This is expected since, for large values of α , it is easier to detect the clients with very noisy data, and thus their contribution to the overall model decreases.

Table 8

Comparison of the RMSE and MAE (in flights) of the RUL prognostics for all three test engines in N-CMAPSS, for all considered aggregation methods. Here, additional noise with a standard deviation multiplier of $\alpha = 1$ is added to the data of engine 5 and 18. The best results are denoted in bold.

Aggregation Method	Engine 11		Engine 14		Engine 15		All engines	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
FedAvg	13.8	11.8	11.9	9.7	11.5	9.5	12.3	10.2
Softmax aggregation policy								
Random-Softmax	7.6	6.3	13.9	9.7	8.0	6.3	10.5	7.6
Full-Softmax	5.9	5.0	13.7	9.5	8.0	5.9	10.1	7.0
Best model aggregation policy								
Random-Best	5.6	4.4	19.1	14.2	10.1	7.1	13.4	9.0
Full-Best	6.1	5.1	12.6	9.1	7.7	5.9	9.5	6.9

For $\alpha = 0.7$ onwards, a lower RMSE is obtained over all test engines for both the Full-Softmax and the Random-Softmax method. Here, the RMSE over all test engines with the Full-Softmax method, is lower and more stable with the Random-Softmax method. This is as expected since the Random-Softmax method has an extra source of randomness, compared to the Full-Softmax method, because local models are randomly assigned to clients. This leads to less robust outcomes.

The outcomes with the best model aggregation policies vary with different level of noise. This is because for both Full-Best and the Random-Best method, the impact of randomness is larger, since the best model is selected as the new global model in each iteration. A small, random change in the training process, for instance slightly different initial weights, might already give another best model with the Full-Best and Random-Best methods.

Table 9 shows how often the parameters (model weights) of each airline are selected as the parameters of the new global model with the Random-Best aggregation method. For a noise multiplier of $\alpha = 0.1$ added to the data of engines 5 and 18, engine 5 is selected 3 times, while engine 18 is selected 17 times out of 50. In contrast, for a higher noise multiplier of $\alpha = 1.0$ added to the data of engines 5 and 18, engine 5 is no longer selected, while engine 18 is only selected 8 times out of 50. This shows that the magnitude of the noise present in an engine's data impacts the frequency at which the engine is selected for the global model, which is in line with the results in Tables 7 and 8. Notably, the engine that performs the best when trained in isolation (Engine 16 in Table 5), is also one of the engines selected most frequently in both noise conditions. This is a desirable outcome, as training a model in isolation with Engine 16's data provided good test results, as shown in Table 5. Overall, Table 9 shows the variety in the contributions of all models trained on different engine datasets, in spite of their dissimilarity to the validation dataset data distributions, promoting the inclusion of features from various clients.

6.4. Limitations of this work

Privacy vs. robustness in the aggregation methods. The proposed aggregation methods enhance robustness to excessive noise in the data of some clients. For these aggregation methods, the server shares the weights of the local model of an airline with the other airlines. However, information about the training data of each airline is stored in its local weights [33]. An airline could therefore perform inversion attacks and property inference attacks to infer information about, and characteristics of, the training data of another airline from the model weights [33,48,49]. There is thus a trade-off between the robustness of the approach and the privacy protection of the airlines.

The impact the number of clients with noisy data on the performance of the proposed aggregation methods. In the context of distributed learning, it is expected that noise may appear in some of the client's (airline's) datasets. Nevertheless, given that the methods presented in this paper

rely on the statistical median's robustness to outliers, as soon as more than half the clients contain noise in their data, the median can no longer be used as a suitable selector. Therefore, these methods are limited to contexts where no more than half the clients will present noise in their datasets. As a result, different approaches should be considered for such cases.

Approval from the regulatory authorities. Prognostic models for aircraft have to be verified and approved by regulatory authorities such as the European Union Aviation Safety Agency (EASA, [22]) in Europe, or the Federal Aviation Administration (FAA) in the USA. For this, it might be necessary that airlines share aggregated information about their data. In general, however, a single airline may not have enough data to get a prognostic model approved by the regulatory authorities. The methodology proposed in this paper is a first attempt to develop a collaborative framework towards getting approval from the regulatory authorities.

Stragglers. Another limitation of the proposed framework is the lack of support to mitigate stragglers in a heterogeneous setup [50], where the speed of the model training is determined by the slowest device. This might hinder the speed of training of the RUL prognostic model. Nevertheless, it is expected that airlines are interested in investing in appropriate hardware and software capabilities to support a scalable implementation of the proposed FL framework.

7. Conclusion

In this paper, a feasible federated learning framework is developed that enables airlines to collaboratively train a common machine learning model for RUL prognostics, without the need to share their data. Validating the common machine learning model for RUL prognostics is crucial for airlines to be able to certify and implement such model in practice. However, to address the reluctance to collaborate due to data sharing and potential noisy data from peers, a decentralized validation procedure and four aggregation methods are proposed within the FL framework.

The proposed FL framework is applied for the development of RUL prognostic models for aircraft engines [23]. Six airlines are considered. Using the proposed FL framework, accurate RUL prognostics are obtained, with a RMSE of 9.9 flights. The performance of the FL framework is compared with the case where airlines do not cooperate, i.e., each airline is independently training its own RUL prognostic model. In this case, the mean RMSE (over all six airlines) is 15.8 flights, showing a significant decrease in the accuracy of the RUL prognostics. The FL approach provides more accurate RUL prognostics for five out of the six considered airlines. Overall, the results show that the proposed FL framework gives accurate RUL prognostics, while the privacy of their data is preserved. Lastly, four novel aggregation methods are proposed for the global model parameters. When all airlines have no or little noise in the data, the proposed methods and the standard FedAvg method have a comparable performance. However, when some airlines have noisy condition monitoring data, the proposed aggregation methods provide more accurate RUL prognostics.

Overall, this paper investigates the robustness property of FL, allowing the construction of effective global models, even in the presence of noise. However, as pointed out in [33,48,49], FL procedures require special attention to protect a client's privacy. As such, our future work considers the development of such procedures and explore privacy-preserving collaborative learning mechanisms, while at the same time remaining robust to noisy data. Such approaches will also consider authentication as an additional measure to improve the robustness of aggregation methods. Nevertheless, this paper has shown how multiple airlines would benefit in collaborating for the development of a global RUL prognostics model. At the same time, these methods ensure a minimal impact on the resulting model from clients with noisy data.

Table 9

An overview of how often the parameters (model weights) of each airline are selected as the parameters of the new global model, and in which epochs these airlines are selected, for the Random-Best aggregation method with $\alpha = 0.1$ (small noise) and $\alpha = 1.0$ (large noise).

Noise multiplier	Airline (Engine)	# Times selected	Selected in epochs
$\alpha = 0.1$	Airline A (Engine 2)	7	2, 3, 19, 20, 27, 40, 49
	Airline B (Engine 5)	3	8, 24, 50
	Airline C (Engine 10)	10	4, 11, 12, 17, 34, 35, 38, 43, 46, 47,
	Airline D (Engine 16)	12	7, 10, 21, 25, 26, 30, 31, 32, 36, 41, 45, 48
	Airline E (Engine 18)	17	1, 5, 6, 9, 13, 14, 15, 16, 18, 22, 23, 28, 29, 33, 37, 39, 42
	Airline F (Engine 20)	1	44
$\alpha = 1.0$	Airline A (Engine 2)	3	8, 13, 35
	Airline B (Engine 5)	0	–
	Airline C (Engine 10)	15	1, 3, 4, 10, 12, 14, 16, 22, 23, 25, 31, 32, 37, 39, 44
	Airline D (Engine 16)	17	9, 11, 15, 19, 26, 27, 28, 30, 33, 34, 36, 38, 41, 42, 45, 47, 48
	Airline E (Engine 18)	8	6, 17, 21, 24, 29, 46, 49, 50
	Airline F (Engine 20)	7	2, 5, 7, 18, 20, 40, 43

CRedit authorship contribution statement

Diogo Landau: Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Ingeborg de Pater:** Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Mihaela Mitici:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Investigation, Conceptualization. **Nishant Saurabh:** Writing – review & editing, Writing – original draft, Supervision, Methodology, Investigation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix. Artifact description

We briefly present in this section the reproducibility artifact of our proposed FL framework.

A.1. Release

The complete prototype of the proposed FL framework and its source code is released open-source [24] in an *online GitHub* repository at <https://github.com/EC-labs/rul-fl>.

A.2. Hardware dependencies

We implemented and tested the proposed FL framework on an i7-11850H Intel server @ 2.50GHz with 8 CPU cores, 2×16GB DDR4 RAM, and an encryption enabled 1TB SSD disk. The software, however, is portable and has no hardware dependencies.

A.3. Software dependencies

We tested the proposed FL framework on an Ubuntu 22.04 LTS, Jammy Jellyfish (x86_64) OS with 6.8.12 Linux kernel, however it runs on any Linux-based environment. The implemented framework has no other software dependencies.

A.4. Dataset

The complete dataset [51] used for illustrating the usefulness of the proposed FL framework is available at <https://pdm-datasets.s3.amazonaws.com/NASA/17.+Turbofan+Engine+Degradation+Simulation+Data+Set+2.zip>. The data preprocessing details are presented in Section 5.1.

A.5. Evaluation data

The complete data used for analysis in the paper is available at *GitHub* repository [24] in the compressed format, and can be found in the `results` directory as:

```
./results/evaluation=2024-03-14.zip
```

Reproducing the analysis data in the `results` directory requires executing `run.sh` script in the *GitHub* repository [24]. Currently, as part of FL configuration, the `run.sh` script runs 5 clients and 1 server for each federated learning algorithm (i.e., fedavg, full-softmax, full-best, random-softmax, random-best) presented in the paper. The `run.sh` script also runs experiments for the unrestricted access centralized, and non-collaborative isolated learning scenarios.

Additionally, the `run.sh` script also runs experiments for multiple noise configurations, starting at $\alpha = 0.1$ to $\alpha = 2.0$. The analysis data related to complete experimentation can be found in the `results/` directory.

Data availability

The data used is mentioned in the paper, and is publicly available.

References

- [1] O. Fink, Q. Wang, M. Svendsen, P. Dersin, W.-J. Lee, M. Ducoffe, Potential, challenges and future directions for deep learning in prognostics and health management applications, *Eng. Appl. Artif. Intell.* 92 (2020) 103678.
- [2] M. Mitici, I. de Pater, A. Barros, Z. Zeng, Dynamic predictive maintenance for multiple components using data-driven probabilistic RUL prognostics: The case of turbofan engines, *Reliab. Eng. Syst. Saf.* 234 (2023) 109199.
- [3] I. de Pater, M. Mitici, Developing health indicators and RUL prognostics for systems with few failure instances and varying operating conditions using a LSTM autoencoder, *Eng. Appl. Artif. Intell.* 117 (2023) 105582.
- [4] W. Zhang, X. Li, H. Ma, Z. Luo, X. Li, Federated learning for machinery fault diagnosis with dynamic validation and self-supervision, *Knowl.-Based Syst.* 213 (2021) 106679.
- [5] SESAR, Final project results report: AICHAIN, a platform for privacy-preserving federated machine learning using blockchain to enable operational improvements in ATM, 2022.
- [6] S. Ruiz, A. Rungger, J. Busto, S. Toor, I. Martin, Privacy-preserving federated machine learning in ATM: experimental results from two use cases, *Proc. 12th SESAR Innov. Days* (2022).
- [7] B. McMahan, E. Moore, D. Ramage, S. Hampson, B.A. y Arcas, Communication-efficient learning of deep networks from decentralized data, in: *Artificial Intelligence and Statistics*, PMLR, 2017, pp. 1273–1282.
- [8] Y. Lei, N. Li, L. Guo, N. Li, T. Yan, J. Lin, Machinery health prognostics: A systematic review from data acquisition to RUL prediction, *Mech. Syst. Signal Process.* 104 (2018) 799–834.

- [9] B. Rezaeianjouybari, Y. Shang, Deep learning for prognostics and health management: State of the art, challenges, and opportunities, *Measurement* 163 (2020) 107929.
- [10] A.N. Bhagoji, S. Chakraborty, P. Mittal, S. Calo, Analyzing federated learning through an adversarial lens, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 634–643.
- [11] L. Guo, Y. Yu, M. Qian, R. Zhang, H. Gao, Z. Cheng, FedRUL: A new federated learning method for edge-cloud collaboration based remaining useful life prediction of machines, *IEEE/ASME Trans. Mechatronics* (2022).
- [12] Y. Wang, T. Zhu, W. Chang, S. Shen, W. Ren, Model poisoning defense on federated learning: A validation based approach, in: *International Conference on Network and System Security*, Springer, 2020, pp. 207–223.
- [13] D. Yin, Y. Chen, R. Kannan, P. Bartlett, Byzantine-robust distributed learning: Towards optimal statistical rates, in: *International Conference on Machine Learning*, PMLR, 2018, pp. 5650–5659.
- [14] P. Blanchard, E.M. El Mhamdi, R. Guerraoui, J. Stainer, Machine learning with adversaries: Byzantine tolerant gradient descent, *Adv. Neural Inf. Process. Syst.* 30 (2017).
- [15] W. Zhang, X. Li, Federated transfer learning for intelligent fault diagnostics using deep adversarial networks with data privacy, *IEEE/ASME Trans. Mechatronics* 27 (1) (2021) 430–439.
- [16] H. Lu, A. Thelen, O. Fink, C. Hu, S. Laflamme, Federated learning with uncertainty-based client clustering for fleet-wide fault diagnosis, 2023, arXiv preprint arXiv:2304.13275.
- [17] J. Chen, J. Li, R. Huang, K. Yue, Z. Chen, W. Li, Federated transfer learning for bearing fault diagnosis with discrepancy-based weighted federated averaging, *IEEE Trans. Instrum. Meas.* 71 (2022) 1–11.
- [18] K. Zhao, J. Hu, H. Shao, J. Hu, Federated multi-source domain adversarial adaptation framework for machinery fault diagnosis with data privacy, *Reliab. Eng. Syst. Saf.* 236 (2023) 109246.
- [19] W. Zhang, X. Li, Data privacy preserving federated transfer learning in machinery fault diagnostics using prior distributions, *Struct. Heal. Monit.* 21 (4) (2022) 1329–1344.
- [20] J. Chen, J. Li, R. Huang, K. Yue, Z. Chen, W. Li, Federated learning for bearing fault diagnosis with dynamic weighted averaging, in: *2021 International Conference on Sensing, Measurement & Data Analytics in the Era of Artificial Intelligence*, ICSMD, IEEE, 2021, pp. 1–6.
- [21] M. Fang, X. Cao, J. Jia, N. Gong, Local model poisoning attacks to {Byzantine-robust} federated learning, in: *29th USENIX Security Symposium*, USENIX Security 20, 2020, pp. 1605–1622.
- [22] EASA, Concept Paper: First Usable Guidance for Level 1 Machine Learning Applications, Technical Report, European Union Aviation Safety Agency, 2023.
- [23] M. Arias Chao, C. Kulkarni, K. Goebel, O. Fink, Aircraft engine run-to-failure dataset under real flight conditions for prognostics and diagnostics, *Data* 6 (1) (2021) 5.
- [24] Federated learning framework for collaborative remaining useful life prognostics: an aircraft engine case study, 2025, <https://github.com/EC-labs/rul-fl>. (Accessed 20 May 2025).
- [25] H. Li, C. Li, J. Wang, A. Yang, Z. Ma, Z. Zhang, D. Hua, Review on security of federated learning and its application in healthcare, *Future Gener. Comput. Syst.* 144 (2023) 271–290.
- [26] Y. Djenouri, T.P. Michalak, J.C.-W. Lin, Federated deep learning for smart city edge-based applications, *Future Gener. Comput. Syst.* 147 (2023) 350–359.
- [27] F.N. Al-Wesabi, H.A. Mengash, R. Marzouk, N. Alruwais, R. Allafi, R. Alabdan, M. Alharbi, D. Gupta, Pelican optimization algorithm with federated learning driven attack detection model in internet of things environment, *Future Gener. Comput. Syst.* (2023).
- [28] B. Mihai, N. Saurabh, D. Kimovski, Multi-objective optimization of federated learning systems in the computing continuum, in: M. Farmanbar, M. Tzamtzi, A.K. Verma, A. Chakravorty (Eds.), *Frontiers of Artificial Intelligence, Ethics, and Multidisciplinary Applications*, Springer Nature Singapore, Singapore, ISBN: 978-981-99-9836-4, 2024, pp. 141–154.
- [29] W. Yang, J. Chen, Z. Chen, Y. Liao, W. Li, Federated transfer learning for bearing fault diagnosis based on averaging shared layers, in: *2021 Global Reliability and Prognostics and Health Management*, PHM-Nanjing, IEEE, 2021, pp. 1–7.
- [30] H. Lu, A. Thelen, O. Fink, C. Hu, S. Laflamme, Federated learning with uncertainty-based client clustering for fleet-wide fault diagnosis, *Mech. Syst. Signal Process.* 210 (2024) 111068.
- [31] T. Li, A.K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, V. Smith, Federated optimization in heterogeneous networks, 2018, arXiv preprint arXiv:1812.06127.
- [32] S. Ji, S. Pan, G. Long, X. Li, J. Jiang, Z. Huang, Learning private neural language modeling with attentive aggregation, in: *2019 International Joint Conference on Neural Networks, IJCNN*, IEEE, 2019, pp. 1–8.
- [33] Z. Alebouyeh, A.J. Bidgoly, Benchmarking robustness and privacy-preserving methods in federated learning, *Future Gener. Comput. Syst.* 155 (2024) 18–38.
- [34] J. Hayes, O. Ohrimenko, Contamination attacks and mitigation in multi-party machine learning, *Adv. Neural Inf. Process. Syst.* 31 (2018).
- [35] G. Damaskinos, R. Guerraoui, R. Patra, M. Taziki, et al., Asynchronous Byzantine machine learning (the case of SGD), in: *International Conference on Machine Learning*, PMLR, 2018, pp. 1145–1154.
- [36] C. Xie, S. Koyejo, I. Gupta, Zeno: Distributed stochastic gradient descent with suspicion-based fault-tolerance, in: *International Conference on Machine Learning*, PMLR, 2019, pp. 6893–6901.
- [37] C. Xie, O. Koyejo, I. Gupta, Generalized byzantine-tolerant sgd, 2018, arXiv preprint arXiv:1802.10116.
- [38] X.-C. Li, D.-C. Zhan, Y. Shao, B. Li, S. Song, Fedphp: Federated personalization with inherited private models, in: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, Springer, 2021, pp. 587–602.
- [39] T. Yu, E. Bagdasaryan, V. Shmatikov, Salvaging federated learning by local adaptation, 2020, arXiv preprint arXiv:2002.04758.
- [40] M. Röder, P. Kowalczyk, F. Thiesse, Tracing down the value of co-creation in federated AI ecosystems, in: *ECIS*, 2022.
- [41] E. Sannara, F. Portet, P. Lalanda, 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.
- [42] N. Fei, Y. Gao, Z. Lu, T. Xiang, Z-score normalization, hubness, and few-shot learning, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 142–151.
- [43] M.A. Chao, C. Kulkarni, K. Goebel, O. Fink, Fusing physics-based and deep learning models for prognostics, *Reliab. Eng. Syst. Saf.* 217 (2022) 107961.
- [44] X. Li, Q. Ding, J.-Q. Sun, Remaining useful life estimation in prognostics using deep convolution neural networks, *Reliab. Eng. Syst. Saf.* 172 (2018) 1–11.
- [45] B. Wang, Y. Lei, N. Li, T. Yan, Deep separable convolutional network for remaining useful life prediction of machinery, *Mech. Syst. Signal Process.* 134 (2019) 106330.
- [46] J. Shang, D. Xu, H. Qiu, L. Gao, C. Jiang, P. Yi, A novel data augmentation framework for remaining useful life estimation with dense convolutional regression network, *J. Manuf. Syst.* 74 (2024) 30–40.
- [47] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, 2014, arXiv preprint arXiv:1412.6980.
- [48] L. Zhu, Z. Liu, S. Han, Deep leakage from gradients, *Adv. Neural Inf. Process. Syst.* 32 (2019).
- [49] J. Geiping, H. Bauermeister, H. Dröge, M. Moeller, Inverting gradients-how easy is it to break privacy in federated learning? *Adv. Neural Inf. Process. Syst.* 33 (2020) 16937–16947.
- [50] D. Wu, R. Ullah, P. Harvey, P. Kilpatrick, I. Spence, B. Varghese, FedAdapt: Adaptive offloading for IoT devices in federated learning, *IEEE Internet Things J.* 9 (21) (2022) 20889–20901, <http://dx.doi.org/10.1109/JIOT.2022.3176469>.
- [51] Turbofan dataset, 2025, <https://pdm-datasets.s3.amazonaws.com/NASA/17.+Turbofan+Engine+Degradation+Simulation+Data+Set+2.zip>. (Accessed 20 May 2025).



Diogo H.M.C. Landau is a PhD candidate at the Department of Information and Computing Sciences, Utrecht University, Netherlands. Previously he received his BSc and MSc degree in Electrical and Computer Engineering from the Faculty of Engineering of the University of Porto, Portugal. His research areas include distributed resource monitoring and provisioning, parallel computing, and systems performance modeling.



Ingeborg de Pater is a PhD candidate at the Faculty of Aerospace Engineering at the Delft University of Technology (the Netherlands). She obtained her masters in Econometrics (track: Operations Research) from the Erasmus University Rotterdam. In her PhD, she works on the development of Remaining Useful Life (RUL) prognostic models, where the uncertainty of the RUL prognostics is also quantified, and on the integration of these uncertain RUL prognostics in the maintenance schedule optimization. She applies this to predictive maintenance for aircraft.



Mihaela Mitici is an Assistant Professor at Faculty of Science, Utrecht University, the Netherlands. She has a PhD in Stochastic Operations Research from University of Twente, the Netherlands, and an MSc in Operations Research from University of Amsterdam. During 2016-2022 she was Assistant professor at Aerospace Faculty, TU Delft. Mihaela specializes in Operations Research, with a focus on stochastic processes, decision-making under uncertainty, machine learning. Her main application domains are predictive maintenance, RUL prognostics, and mobility.



Nishant Saurabh is a tenured Assistant Professor in the Department of Information and Computing Sciences at Utrecht University in the Netherlands. He obtained his Ph.D. in Computer Science from the University of Innsbruck in 2021 and later worked as a postdoctoral researcher at Klagenfurt University, Austria. His research interests include hybrid distributed systems, cloud and edge computing, performance modeling, optimization, and observability. He is an associate editor for Springer's JoCCASA journal, editorial board and steering committee member for Springer's book series and conference on frontiers of AI. He also served as scientific coordinator and WP leader in several EU and Austrian projects and is currently a member of IBM's working committee on HPC-Quantum integration.