



WeightRelay: Efficient Heterogeneous Federated Learning on Time Series

Wensi Tang and Guodong Long^(✉)

Australian Artificial Intelligence, Institute Faculty of Engineering and IT,
University of Technology Sydney, Ultimo, Australia
wensi.tang@student.uts.edu.au, guodong.long@uts.edu.au

Abstract. Federated learning for heterogeneous devices aims to obtain models of various structural configurations in order to fit multiple devices according to their hardware configurations and external environments. Existing solutions train those heterogeneous models simultaneously, which requires extra cost (e.g. computation, communication, or data) to transfer knowledge between models. In this paper, we proposed a method, namely, weight relay (WeightRelay), that could get heterogeneous models without any extra training cost. Specifically, we find that, compared with the classic random weight initialization, initializing the weight of a large neural network with the weight of a well-trained small network could reduce the training epoch and still maintain a similar performance. Therefore, we could order models from the smallest and train them one by one. Each model (except the first one) can be initialized with the prior model's trained weight for training cost reduction. In the experiment, we evaluate the weight relay on 128-time series datasets from multiple domains, and the result confirms the effectiveness of WeightRelay. More theoretical analysis and code can be found in (https://github.com/Wensi-Tang/DPSN/blob/master/AJCAI23_wensi_fedTSC.pdf).

Keywords: Time series classification · Federated learning · Heterogeneous model

1 Introduction

With the development of smart devices, an increasing amount of time series data can be collected, such as daily heartbeats, blood oxygen levels, electronic consumption, and motion signals for smart device control [1, 6–8, 19, 21, 35].

With the advent of sophisticated devices and advanced data analysis technologies, there is potential to bring immense value to society. However, privacy concerns limit the integration of smart devices with state-of-the-art deep learning [24, 59]. Specifically, smart devices encounter challenges when training large models due to limited power and computational resources for local training. Concurrently, uploading information from these devices, which are commonly found in homes, can pose significant privacy risks, as they often contain sensitive data [14, 23].

Federated learning [14] can be used to train deep learning models with privacy protection. However, there are still many unsolved challenges in applying those solutions to smart devices. Specifically, smart devices even for a similar function, typically have varying hardware configurations. The hardware heterogeneous brings a challenge for federated learning. Specifically, how to transfer knowledge between heterogeneous models [14, 21, 59]. Specifically, this heterogeneous federated learning aims to obtain models of various structural configurations to fit multiple devices according to their hardware configurations and working environment. Under this setting, it is hard for low-capacity devices to contribute their knowledge to big models, for they might have enough memory, bandwidth or computational power to join the big model training via Federate average. Therefore, solutions that could enable big models to get knowledge from the small models are highly desired [14, 21].

Existing solutions tackle the problem by adding one or more resources such as computational, communication, and extra data. For example, distillation-based methods require training cost on multiple models and extra cost for knowledge transfer between those models [20, 30, 31, 58, 61]. Pruning-based methods [25, 29] require an extra cost to pruning the single model onto multiple smaller models. Weight sharing methods, some of them [51] require the computational cost to match the weight of various models iteratively, and some methods [11, 55] need a weight scale module to adjust weight before sharing.

Although existing solutions enable knowledge sharing between heterogeneous models, the extra resource consumption makes it hard to implement them on smart devices. This is because most of those devices do not have strong computational, memory or commutations capacity [14, 21]. Therefore, solutions with huge training costs bring an embarrassing burden to low-capacity smart devices. When a device's ability limits it from using big models, it may also limit it from contributing its knowledge to big models or getting knowledge from big models via distillations, pruning or weight matching.

Other than capacity neglection, for smart devices, the appropriate 1D-CNN (one-dimensional-convolutional neural networks) model is also seldom mentioned. Specifically, most solutions were tested with 2D-CNN(two-dimensional-convolutional neural networks) models [11, 20, 30, 51, 55, 58, 61] or language models [31]. However, we should notice that most of the data gathered by smart devices are time-series data [33, 53] which can mathematically be described as a series of data points recorded in time order [8, 9, 13, 49]. Such as the heartbeat data collected by smartwatches [37, 39], the electricity consumption data gathered by energy management devices [15], building structural vibration data recorded by motion sensors [27, 50], etc. According to the University of California, Riverside time series archive (UCR archive) [9], the state-of-the-art solutions for time series classification tasks are all 1D-CNNs [10, 13, 49].

The characteristics of 1D-CNN allow a novel weight relay solution, which does not need any extra resources. Specifically, suppose we have a small 1D-CNN network and a large 1D-CNN network. For the large network training, we could initialize it with a classic random weight initialization or we could initialize it with the weight from a well-trained small network. We find that these two kinds of initialization will be of similar performance, but the second initialization could reduce the training cost of the large network. This training cost reduction could be used to lower the capacity requirement

and allows more low-capacity devices to join the big model training. Via ordering those heterogeneous 1D-CNN models from the smallest to the largest, except for the first smallest model, all the other models' training will be benefited.

In experiments, we show the consistently training cost reduction ability of the weight relay on time series datasets from multiple domains i.e., healthcare, human activity recognition, speech recognition, and material analysis. Despite the dynamic patterns of these datasets, weight relay robustly shows its effectiveness.

2 Related Work

2.1 Deep Learning for Time Series Classification

The success of deep learning encourages the exploration of its application on time series data [12, 13, 28]. Intuitively, the Recurrent Neural Network (RNN), which is designed for temporal sequence, should work on the time series tasks. However, in practice, RNN is rarely applied to TS classification [13]. One widely accepted reason among many is that RNN models suffer from vanishing and exploding gradients when dealing with long sequence data [2, 13, 38]. Nowadays, 1D-CNN is the most popular deep-learning method TSC tasks. [22, 26, 40, 41, 52, 60]. According to the University of California, Riverside time series archive (UCR archive) [9], the state-of-the-art solutions for time series classification tasks are all 1D-CNNs [10, 22, 49, 52].

2.2 Federated Learning on Heterogeneous Devices

Based on the method of transferring knowledge between heterogeneous models, the solutions to heterogeneous Federated learning could be divided into three columns. These are distillation-based, weight-sharing-based pruning-based methods and prototype-based methods.

The knowledge distillation [5, 18] allows the knowledge sharing between heterogeneous models. It requires extra data or computational resources to enable knowledge sharing between models [20, 30, 56], which brings multiple challenges under the federated learning setting. For example, the distillation process requires a large amount of computational resources [61]. What's more, the performance of the distillation is highly related to the similarity of the distributions between the training data and distillation data.

The weight sharing method is based on the assumption that some parts of the weight of various structure models are the shareable, and the shareable part could help to transfer the knowledge between various structure models [4, 36, 42, 51]. In practice, finding which parts of various models should be of the same parameter is hard. Therefore, a large number of computation resources have to be taken to find which parts of models should be matched together [51], or taken to calculate the adjustment module for weight re-scale [11, 43].

The pruning-based method aims at training a large neural network and pruning it into various small networks [57] according to the hardware configuration and external environment [25, 32, 54]. One limitation is that it is hard to control the structure of

the pruned network, which challenges fitting those small networks according to the configuration of each edge device. [25, 29, 47]

The prototype-based federated learning [3, 16, 17, 46] can also be viewed as a solution for heterogeneous devices for it requires a very limited resource for classification calculation. However, as the number of classes increases, prototype learning struggles to scale effectively. Put simply, due to the warp characteristic, comparing the distance between two-time series incurs a computational cost of N^2 . Where N is the length of the signal. As the number of classes grows, this cost rises substantially. As a result, the computational cost shifts from training to classification, which is not conducive for small devices. While some methods can map the time series into a feature space [45, 48], their classification accuracy cannot match that of larger models.

3 Motivation

To train a 1D-CNN model, we could 1) start from a classic random weight initialization or 2) replace parts of the random initialization weight with a well-trained weight from a small network. The second initialization could reduce the training cost of the large network and won't influence the final performance.

Therefore, to train multiple models on a smaller budget, we don't need to train every model from the stretch. We could initialize some of those models by the trained weight from the others for fast convergence. Figure 1 gives an example of the weight relay on the Crop [44] dataset.




Fig. 1. The left image shows the relationship between the accumulated computational cost and the test accuracy of each model. The accumulated computational cost calculates the computational cost we used to obtain a well-trained model. As the image shows, to obtain a single largest model from classic initialization (purple), we need about $0.8e^7$ computational resources, and we could only get 1 model. However, with weight relay, at the point $0.8e^7$, we have four models (blue, orange, green and red). And the red model, which has the same structure as the purple model, also has the same performance. The right image shows the performance of each model by communication round. We could see that the weight relay model (red) converged much faster than the classic initialization models (purple). (Color figure online)

4 Weight Relay

In Fig. 2, a schematic of weight relay is given. What’s more, in this Section, we will explain the weight relay in detail. Specifically, Sect. 4.1 will introduce heterogeneous models for time series classifications. Section 4.2 will introduce how to align those models. Specifically, when a well-trained weight is passed to a large model, which part of the large network should be replaced.




Fig. 2. The schematic shows training multiple models with weight relay. Weight relay starts from the training of the smallest network with classic initialization. The trained weight of the smallest network will be used to replace a part of the classic initialization weight of a large network. When the large network is trained, its weight could be used to accelerate the training of a larger network. Since the weight relay only replaces the random initialization with a well-trained one, it requires almost no cost, and the training cost of each model (except the first one) is also smaller than training those from classic initialization.

4.1 Heterogeneous Models

According to the result statistics on the UCR archive, all state-of-the-art neural network solutions on time series classification tasks are 1D-CNN models [13, 34, 49, 52]. Therefore, this paper mainly talked about 1D-CNN models and heterogeneous could happen on all three main structure configurations: the number of layers, kernel sizes and the number of channels.

4.2 Weight Alignment

The weight alignment defines which part of the large network should be replaced with the weight of the small network. The alignment will have three steps because the neural network has three hierarchies. Specifically, a neural network is composed of layers. Layers are composed of weight sets. And the weight sets are composed of weight tensors. Therefore, we need to pair weight sets and layers before we align weight tensors.

First Step: Pair Weight Sets by Layers. The convolutional layer and the batch normalization layer will be indexed from the input to the output. For example, the first convolutional layer is the convolutional layer closest to the input. The fully connected layer will be indexed by **reverse order**. Therefore, the first fully connected layer will be the layer closest to the output. For each of those three types of layers, the $*$ -th layer’s tensor set of the small network should be paired with the $*$ -th layer’s tensor set of the large network

Second Step: Pair Tensors by Paired Weight Sets: According to the definitions of the 1D-CNN, tensors in each set will have different functions, such as weight, bias, and running mean. Therefore, for two paired sets, tensors in the two sets will be paired by their function name.

Third Step: Align Two Paired Tensors. For two paired tensors \mathbf{A} and \mathbf{B} , the alignment of the two tensors is to **align the output and input dimensions with the left margin** and **align the kernel dimension with the centre**. Specifically, the small network's $*$ -th element should be aligned with the middle of the large network's $*$ -th element on the input(output) channel. Therefore, the small network's i -th element should be aligned with the middle of the large network's j -th element on the kernel dimension. Then the i and j describe the alignment relationship of \mathbf{A} and \mathbf{B} in the Eq. 1

$$j = i + \lfloor (b - 1)/2 \rfloor - \lfloor (a - 1)/2 \rfloor \quad (1)$$

where the b and a are kernel sizes of large network and small network.

5 Analysis of Weight Relay

This section will have two parts. In Sect. 5.1, we will show that though the alignment method in Eq. 1 is defined in pair, it is reliable for the alignment of multiple models. Secondly, we will give a macro (Sect. 5.2) and a micro (Sect. 5.3) explanation of the coverage acceleration of the weight relay.

5.1 Consistency Proof for the Alignment

This section will show that despite we only define the relationship between two kernels, this operation will keep the **consistency** when we have **multiple** kernels. The consistency of weight alignment can be describe as below: For any three kernel weights $\{\mathbf{A}, \mathbf{B}, \mathbf{C}\}$ and their length relationship are:

$$a < b < c$$

We could use weight alignment to deter the alignment relationship between \mathbf{C} and the other two kernels as:

$$\mathbf{A}_i \xrightarrow{\text{align with}} \mathbf{C}_k \quad (2)$$

$$\mathbf{B}_j \xrightarrow{\text{align with}} \mathbf{C}_k \quad (3)$$

The signal Phase alignment can be called consistency if, with the same operation on \mathbf{A} and \mathbf{B} , we should have

$$\mathbf{A}_i \xrightarrow{\text{align with}} \mathbf{B}_j$$

Here, we give an example to illustrate the consistency of multiple kernels. Supposing we have three kernels of length 3, 5, 8. Via Eq. 1, the \mathbf{A}_1 should align with \mathbf{C}_4 , the \mathbf{B}_3 should align with \mathbf{C}_4 . **Consistency** means that the \mathbf{A}_1 should align with \mathbf{B}_3 .

Proof of Consistency for Eq. 1:

When we combine Eq. 1 and Eq. 2 we know that the a th element of kernel \mathbf{A} should align with c th element of kernel \mathbf{C} and their index relationship is:

$$k = i + \lfloor (c - 1) / 2 \rfloor - \lfloor (a - 1) / 2 \rfloor \quad (4)$$

With Eq. 1 and Eq. 3 we know the alignment relationship between elements in \mathbf{B} and \mathbf{C} is:

$$k = j + \lfloor (c - 1) / 2 \rfloor - \lfloor (b - 1) / 2 \rfloor \quad (5)$$

Using the Eq. 5 to subtract Eq. 4 and we have:

$$k - k = \begin{matrix} i + \lfloor (c - 1) / 2 \rfloor - \lfloor (a - 1) / 2 \rfloor \\ + \\ -j - \lfloor (c - 1) / 2 \rfloor + \lfloor (b - 1) / 2 \rfloor \end{matrix} \quad (6)$$

Therefore, we know

$$j \equiv i + \lfloor (b - 1) / 2 \rfloor - \lfloor (a - 1) / 2 \rfloor$$

Which means that

$$\mathbf{A}_i \xrightarrow{\text{should align with}} \mathbf{B}_j$$

as it should be.




Fig. 3. This example shows that: Given the same input and target, the gradient direction of the small kernel is the same as the gradient direction of the aligned part of the large network. Specifically, images from left to right are: 1) input and target, which are generated by random noise; 2) A large kernel weight which is random noise, and the weight of the small kernel, which is cropped from the large kernel; 3) The convolution results for large kernel and small kernel with the input in image 1); 4) The gradients of kernels which is calculated with the random target in image 1). To demonstrate the overlapping parts of the two kernels are in a similar direction, their magnitude was adjusted; 5). A zoom view of the overlapping gradient parts (indexed 70 to 90).

5.2 Macro Explanation of the Training Acceleration

One explanation of the training acceleration is that compared with the classic initialization weight, the weight relay initialization is closer to the final weight. In Fig. 4 We statistic the distance between the initialization weight and trained weight, and we will see that compared with the random initialization, the weight relay initialization are of a smaller distance to the final value.




Fig. 4. The blue points are results from 5000 tests with random input, targets, and weight. The orange point is the origin point. We could see that when the value of the x-axis approaches 0 (target similar), the value of the y-axis comes to 0 (gradient direction similar). (Color figure online)

5.3 Micro Explanation of the Training Acceleration

The micro explanation to the training acceleration is that: the small network's optimization direction is the optimization direction of the large network's sub-network, where the small network should replace the weight. Therefore, the training acceleration is because the training of the small network makes the sub-network close to the final target.

To explain this, we could start from a simple case study in Fig. 3.

6 Experiment

6.1 Benchmarks

University of California, Riverside (UCR) 128 archive [9] is selected to evaluate the weight relay under the federated setting. This is an archive of 128 univariate TS datasets from various domains, such as speech reorganizations, health monitoring, and spectrum analysis. What's more, those datasets also have different characteristics. For instance, among those datasets, the class number varies from 2 to 60, and the length of each dataset varies from 24 to 2844. The number of training data varies from 16 to 8,926.

6.2 Evaluation Criteria

Following the datasets archive [9], the accuracy score is selected to measure the performance. Following the paper, the communication round multiple model size is selected to measure the training cost.

6.3 Experiment Setup

Following [13,22,49], for all benchmarks, we follow the standard and unify settings [52] for all 128 datasets in the UCR archive. The training will stop when the training loss is less than $1e-3$ or reach 5000 epoch. To mimic the federated learning scenario, the client number is 10. More details can be found in the supplementary material.

6.4 Experiment Result

The experiment shows that weight relay has similar performance and fewer computation resources costs than using federated average on all devices to obtain each model. Due to the larger number of datasets, we cannot list the results of all datasets. Therefore, we plot the statistical result of the 128 datasets, and the result is shown in Fig. 5.

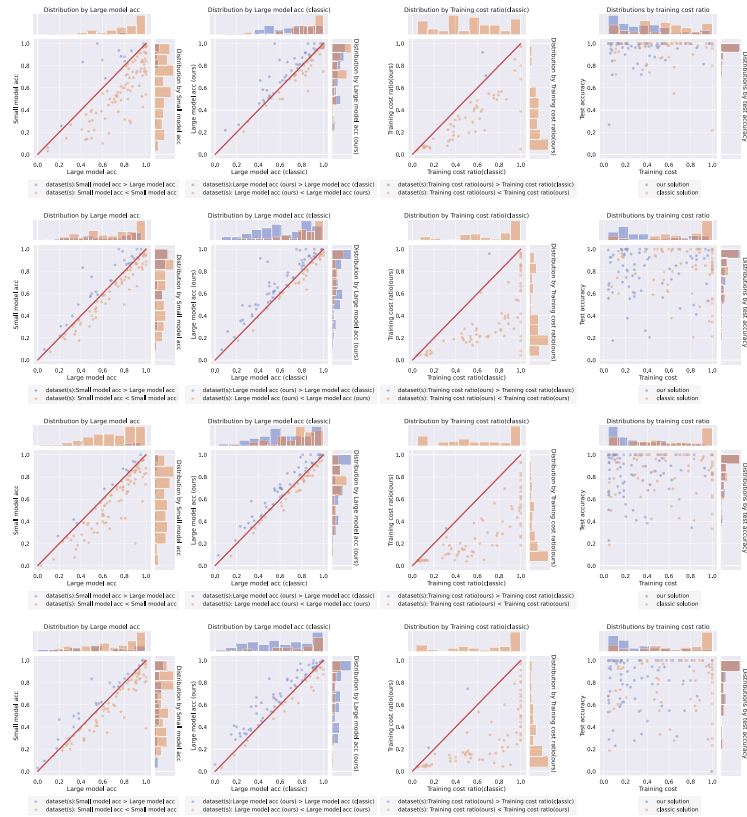


Fig. 5. From top to bottom, each row shows the statistical results when the large model has 2Xkernel size, 2Xchannel number, one extra layer, and all extensions than the smaller model. From the left to the right image, we can see that, for most of the dataset, the large model will perform better than the small model (the first column). The performance of the weight relay is similar to the performance of training from classic (the second column); the weight relay has a lower training cost (the third column); To archive similar performance, the weight relay has a smaller training cost (the fourth column).

7 Conclusion

In this paper, we proposed the weight relay method, which could reduce the training cost for heterogeneous model training. We theoretically analyse the mechanism of weight relay and experimentally verify the effectiveness on multiple datasets from multiple domains.

Acknowledgements. Please place your acknowledgments at the end of the paper, preceded by an unnumbered run-in heading (i.e. 3rd-level heading).

References

1. Bagnall, A., et al.: The UEA multivariate time series classification archive, 2018. arXiv preprint [arXiv:1811.00075](https://arxiv.org/abs/1811.00075) (2018)
2. Bengio, Y., Simard, P., Frasconi, P.: Learning long-term dependencies with gradient descent is difficult. *IEEE Trans. Neural Netw.* **5**(2), 157–166 (1994)
3. Biehl, M., Hammer, B., Villmann, T.: Prototype-based models in machine learning. *Wiley Interdisc. Rev.: Cogn. Sci.* **7**(2), 92–111 (2016)
4. Cai, H., Gan, C., Wang, T., Zhang, Z., Han, S.: Once-for-all: train one network and specialize it for efficient deployment. arXiv preprint [arXiv:1908.09791](https://arxiv.org/abs/1908.09791) (2019)
5. Chen, F., Long, G., Wu, Z., Zhou, T., Jiang, J.: Personalized federated learning with graph. arXiv preprint [arXiv:2203.00829](https://arxiv.org/abs/2203.00829) (2022)
6. Chen, S., Long, G., Shen, T., Jiang, J.: Prompt federated learning for weather forecasting: toward foundation models on meteorological data. arXiv preprint [arXiv:2301.09152](https://arxiv.org/abs/2301.09152) (2023)
7. Chen, S., Long, G., Shen, T., Zhou, T., Jiang, J.: Spatial-temporal prompt learning for federated weather forecasting. arXiv preprint [arXiv:2305.14244](https://arxiv.org/abs/2305.14244) (2023)
8. Chen, Y., et al.: The UCR time series classification archive (2015). www.cs.ucr.edu/eamonn/time_series_data/
9. Dau, H.A., Bagnall, A., Kamgar, K., et al.: The UCR time series archive. [arXiv:1810.07758](https://arxiv.org/abs/1810.07758) (2018)
10. Dempster, A., Petitjean, F., Webb, G.I.: ROCKET: exceptionally fast and accurate time series classification using random convolutional kernels. *Data Min. Knowl. Disc.* **34**(5), 1454–1495 (2020)
11. Diao, E., Ding, J., Tarokh, V.: Heterofl: computation and communication efficient federated learning for heterogeneous clients. arXiv preprint [arXiv:2010.01264](https://arxiv.org/abs/2010.01264) (2020)
12. Dong, X., Kedziora, D., Musial, K., Gabrys, B.: Automated deep learning: Neural architecture search is not the end. arXiv preprint [arXiv:2112.09245](https://arxiv.org/abs/2112.09245) (2021)
13. Fawaz, H.I., Forestier, G., Weber, J., Idoumghar, L., Muller, P.A.: Deep learning for time series classification: a review. *Data Min. Knowl. Disc.* **33**(4), 917–963 (2019)
14. Ferrag, M.A., Friha, O., Maglaras, L., Janicke, H., Shu, L.: Federated deep learning for cyber security in the internet of things: concepts, applications, and experimental analysis. *IEEE Access* **9**, 138509–138542 (2021)
15. Gans, W., Alberini, A., Longo, A.: Smart meter devices and the effect of feedback on residential electricity consumption: evidence from a natural experiment in northern ireland. *Energy Econ.* **36**, 729–743 (2013)
16. Gee, A.H., Garcia-Olano, D., Ghosh, J., Paydarfar, D.: Explaining deep classification of time-series data with learned prototypes. In: *CEUR Workshop Proceedings*, vol. 2429, p. 15. NIH Public Access (2019)
17. Ghods, A., Cook, D.J.: PIP: pictorial interpretable prototype learning for time series classification. *IEEE Comput. Intell. Mag.* **17**(1), 34–45 (2022)
18. Gou, J., Yu, B., Maybank, S.J., Tao, D.: Knowledge distillation: a survey. *Int. J. Comput. Vision* **129**(6), 1789–1819 (2021)
19. Gu, P., et al.: Multi-head self-attention model for classification of temporal lobe epilepsy subtypes. *Front. Physiol.* **11**, 1478 (2020)
20. He, C., Annaram, M., Avestimehr, S.: Group knowledge transfer: federated learning of large CNNs at the edge. *Adv. Neural. Inf. Process. Syst.* **33**, 14068–14080 (2020)
21. Imteaj, A., Thakker, U., Wang, S., Li, J., Amimi, M.H.: A survey on federated learning for resource-constrained IoT devices. *IEEE Internet Things J.* **9**(1), 1–24 (2021)
22. Ismail Fawaz, H., et al.: InceptionTime: finding AlexNet for time series classification. arXiv e-prints [arXiv:1909.04939](https://arxiv.org/abs/1909.04939) (2019)

23. Ji, S., Long, G., Pan, S., Zhu, T., Jiang, J., Wang, S.: Detecting suicidal ideation with data protection in online communities. In: Li, G., Yang, J., Gama, J., Natwchai, J., Tong, Y. (eds.) DASFAA 2019. LNCS, vol. 11448, pp. 225–229. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-18590-9_17
24. Jiang, J., Ji, S., Long, G.: Decentralized knowledge acquisition for mobile internet applications. *World Wide Web* **23**(5), 2653–2669 (2020)
25. Jiang, Y., et al.: Model pruning enables efficient federated learning on edge devices. *IEEE Trans. Neural Netw. Learn. Syst.* (2022)
26. Kashiparekh, K., Narwariya, J., Malhotra, P., Vig, L., Shroff, G.: ConvtimeNet: a pre-trained deep convolutional neural network for time series classification. [arXiv:1904.12546](https://arxiv.org/abs/1904.12546) (2019)
27. Kavyashree, B., Patil, S., Rao, V.S.: Review on vibration control in tall buildings: from the perspective of devices and applications. *Int. J. Dyn. Control* **9**(3), 1316–1331 (2021)
28. Långkvist, M., Karlsson, L., Loutfi, A.: A review of unsupervised feature learning and deep learning for time-series modeling. *Pattern Recogn. Lett.* **42**, 11–24 (2014)
29. Li, A., Sun, J., Li, P., Pu, Y., Li, H., Chen, Y.: Hermes: an efficient federated learning framework for heterogeneous mobile clients. In: Proceedings of the 27th Annual International Conference on Mobile Computing and Networking, pp. 420–437 (2021)
30. Li, D., Wang, J.: Fedmd: Heterogeneous federated learning via model distillation. *arXiv preprint* [arXiv:1910.03581](https://arxiv.org/abs/1910.03581) (2019)
31. Liu, R., et al.: No one left behind: inclusive federated learning over heterogeneous devices. *arXiv preprint* [arXiv:2202.08036](https://arxiv.org/abs/2202.08036) (2022)
32. Liu, S., Yu, G., Yin, R., Yuan, J.: Adaptive network pruning for wireless federated learning. *IEEE Wirel. Commun. Lett.* **10**(7), 1572–1576 (2021)
33. Liu, Y., et al.: Deep anomaly detection for time-series data in industrial IoT: a communication-efficient on-device federated learning approach. *IEEE Internet Things J.* **8**(8), 6348–6358 (2020)
34. Long, G., Shen, T., Tan, Y., Gerrard, L., Clarke, A., Jiang, J.: Federated learning for privacy-preserving open innovation future on digital health. In: Chen, F., Zhou, J. (eds.) *Humanity Driven AI*, pp. 113–133. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-72188-6_6
35. Long, G., Tan, Y., Jiang, J., Zhang, C.: Federated learning for open banking. In: Yang, Q., Fan, L., Yu, H. (eds.) *Federated Learning*. LNCS (LNAI), vol. 12500, pp. 240–254. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-63076-8_17
36. Long, G., Xie, M., Shen, T., Zhou, T., Wang, X., Jiang, J.: Multi-center federated learning: clients clustering for better personalization. *World Wide Web* **26**(1), 481–500 (2023)
37. Park, S., Constantinides, M., Aiello, L.M., Quercia, D., Van Gent, P.: Wellbeat: a framework for tracking daily well-being using smartwatches. *IEEE Internet Comput.* **24**(5), 10–17 (2020)
38. Pascanu, R., Mikolov, T., Bengio, Y.: On the difficulty of training recurrent neural networks. In: *International Conference on Machine Learning*, pp. 1310–1318 (2013)
39. Progonov, D., Sokol, O.: Heartbeat-based authentication on smartwatches in various usage contexts. In: Saracino, A., Mori, P. (eds.) *ETAA 2021*. LNCS, vol. 13136, pp. 33–49. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93747-8_3
40. Rajpurkar, P., Hannun, A.Y., Haghpanahi, M., Bourn, C., Ng, A.Y.: Cardiologist-level arrhythmia detection with convolutional neural networks. [arXiv:1707.01836](https://arxiv.org/abs/1707.01836) (2017)
41. Serrà, J., Pascual, S., Karatzoglou, A.: Towards a universal neural network encoder for time series. In: *CCIA*, pp. 120–129 (2018)
42. Singh, A., Vepakomma, P., Gupta, O., Raskar, R.: Detailed comparison of communication efficiency of split learning and federated learning. *arXiv preprint* [arXiv:1909.09145](https://arxiv.org/abs/1909.09145) (2019)
43. Tan, A.Z., Yu, H., Cui, L., Yang, Q.: Towards personalized federated learning. *IEEE Trans. Neural Netw. Learn. Syst.* (2022)

44. Tan, C.W., Webb, G.I., Petitjean, F.: Indexing and classifying gigabytes of time series under time warping. In: Proceedings of the 2017 SIAM International Conference on Data Mining, pp. 282–290. SIAM (2017)
45. Tan, Y., Liu, Y., Long, G., Jiang, J., Lu, Q., Zhang, C.: Federated learning on non-IID graphs via structural knowledge sharing. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 37, pp. 9953–9961 (2023)
46. Tan, Y., et al.: Fedproto: federated prototype learning across heterogeneous clients. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 36, pp. 8432–8440 (2022)
47. Tan, Y., Long, G., Ma, J., Liu, L., Zhou, T., Jiang, J.: Federated learning from pre-trained models: a contrastive learning approach. *Adv. Neural. Inf. Process. Syst.* **35**, 19332–19344 (2022)
48. Tang, W., Liu, L., Long, G.: Interpretable time-series classification on few-shot samples. In: 2020 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. IEEE (2020)
49. Tang, W., Long, G., Liu, L., Zhou, T., Blumenstein, M., Jiang, J.: Omni-scale CNNs: a simple and effective kernel size configuration for time series classification. In: International Conference on Learning Representations (2021)
50. Vidal, F., Navarro, M., Aranda, C., Enomoto, T.: Changes in dynamic characteristics of Lorca RC buildings from pre-and post-earthquake ambient vibration data. *Bull. Earthq. Eng.* **12**(5), 2095–2110 (2014)
51. Wang, H., Yurochkin, M., Sun, Y., Papailiopoulos, D., Khazaeni, Y.: Federated learning with matched averaging. arXiv preprint [arXiv:2002.06440](https://arxiv.org/abs/2002.06440) (2020)
52. Wang, Z., Yan, W., Oates, T.: Time series classification from scratch with deep neural networks: a strong baseline. In: 2017 International Joint Conference on Neural Networks, pp. 1578–1585. IEEE (2017)
53. Xing, L.: Reliability in internet of things: current status and future perspectives. *IEEE Internet Things J.* **7**(8), 6704–6721 (2020)
54. Xu, W., Fang, W., Ding, Y., Zou, M., Xiong, N.: Accelerating federated learning for IoT in big data analytics with pruning, quantization and selective updating. *IEEE Access* **9**, 38457–38466 (2021)
55. Xu, Z., Yang, Z., Xiong, J., Yang, J., Chen, X.: Elfish: resource-aware federated learning on heterogeneous edge devices. *Ratio* **2**(r1), r2 (2019)
56. Yan, P., Long, G.: Personalization disentanglement for federated learning. arXiv preprint [arXiv:2306.03570](https://arxiv.org/abs/2306.03570) (2023)
57. Zhang, C., et al.: Dual personalization on federated recommendation. arXiv preprint [arXiv:2301.08143](https://arxiv.org/abs/2301.08143) (2023)
58. Zhang, L., Yuan, X.: Fedzkt: zero-shot knowledge transfer towards heterogeneous on-device models in federated learning. arXiv preprint [arXiv:2109.03775](https://arxiv.org/abs/2109.03775) (2021)
59. Zhang, T., Gao, L., He, C., Zhang, M., Krishnamachari, B., Avestimehr, A.S.: Federated learning for the internet of things: applications, challenges, and opportunities. *IEEE Internet Things Mag.* **5**(1), 24–29 (2022)
60. Zheng, Y., Liu, Q., Chen, E., Ge, Y., Zhao, J.L.: Time series classification using multi-channels deep convolutional neural networks. In: Li, F., Li, G., Hwang, S., Yao, B., Zhang, Z. (eds.) *WAIM 2014. LNCS*, vol. 8485, pp. 298–310. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08010-9_33
61. Zhu, Z., Hong, J., Zhou, J.: Data-free knowledge distillation for heterogeneous federated learning. In: International Conference on Machine Learning, pp. 12878–12889. PMLR (2021)