



From federated learning to federated neural architecture search: a survey

Hangyu Zhu¹ · Haoyu Zhang^{2,3} · Yaochu Jin^{1,3}

Received: 11 September 2020 / Accepted: 1 December 2020 / Published online: 4 January 2021
© The Author(s) 2021

Abstract

Federated learning is a recently proposed distributed machine learning paradigm for privacy preservation, which has found a wide range of applications where data privacy is of primary concern. Meanwhile, neural architecture search has become very popular in deep learning for automatically tuning the architecture and hyperparameters of deep neural networks. While both federated learning and neural architecture search are faced with many open challenges, searching for optimized neural architectures in the federated learning framework is particularly demanding. This survey paper starts with a brief introduction to federated learning, including both horizontal, vertical, and hybrid federated learning. Then neural architecture search approaches based on reinforcement learning, evolutionary algorithms and gradient-based are presented. This is followed by a description of federated neural architecture search that has recently been proposed, which is categorized into online and offline implementations, and single- and multi-objective search approaches. Finally, remaining open research questions are outlined and promising research topics are suggested.

Keywords Federated learning · Deep learning · Privacy preservation · Neural architecture search · Reinforcement learning · Evolutionary algorithm · Real-time optimization

Introduction

Deep neural networks (DNNs) have made great success in the fields of image classification, natural language processing, autonomous driving systems, and many others. However, designing DNNs with high-quality architectures usually requires to manually try a large number of different hyperparameters, which is always a tedious task requiring broad expertise in both machine learning and the application area. Therefore, neural architecture search (NAS) has

become increasingly popular in recent years [1], which aims to automatically search for good neural architectures.

Conventional centralized learning systems, however, require that all training data generated on different devices be uploaded to a server or cloud for training a global model, which may give rise to serious privacy concerns. To address this concern, federated learning [2] has been proposed to protect user's data privacy by communicating the model parameters or other model information instead of the raw data between the server and local devices. Naturally, performing NAS in a federated learning environment becomes of particular importance, although it is still in its infant stage.

This survey aims to provide an overview of research work both on federated learning and neural architecture search, focusing, however, on the emerging area of federated neural architecture search. We categorize federated learning systems into offline and online approaches, where online federated neural architecture search is more challenging due to additional requirements on the performance of the networks during the search process and stronger constraints on the computational resources. In addition, we briefly discuss the differences between single- and multi-objective search neural architecture search methods to highlight different ways

✉ Yaochu Jin
yaochu.jin@surrey.ac.uk

Hangyu Zhu
hangyu.zhu@surrey.ac.uk

Haoyu Zhang
zhy920816@sina.cn

¹ Department of Computer Science, University of Surrey, Guildford GU2 7XH, UK

² Engineering Research Center of Digitized Textile & Apparel Technology, Ministry of Education, Shanghai 201620, China

³ College of Information Science and Technology, Donghua University, Shanghai 201620, China

of handling multiple objectives in federated learning, such as accuracy, communication costs, model complexity and memory requirements on the local devices. Finally, we outline the main remaining challenges in federated neural architecture search.

Federated learning

Federated learning [3] distinguishes itself from distributed learning in three aspects. First, the main purpose of federated learning is to protect user's private information while distributed learning aims to accelerate training speed. Second, federated learning cannot determine the data distribution of any client devices. By contrast, distributed learning is able to arbitrarily allocate subsets of the whole learning data. Finally, federated learning faces a more challenging training environment as it may contain millions of unbalanced participating clients whose connections to the server are probably unstable. For example, edge devices like mobile phones are frequently offline.

Federated learning is often categorized based on the distribution characteristics of the data [2] which is originally used in distributed learning. Strictly speaking, federated learning does not have the concept of 'the whole dataset'; therefore, it is hard to accurately describe the federated data distribution to some extent as defined in distribute learning. In the following, we will discuss the data distributions in federated learning in greater detail.

Horizontal federated learning

Horizontal federated learning is proposed for the scenarios in which datasets on the participating clients share the same feature space but have different samples. The name 'horizontal' originates from instance distributed learning [4] as shown in Fig. 1a, where the whole dataset is horizontally partitioned over data samples and allocated to two clients. Similarly, as indicated by the part surrounded by the two dashed lines in Fig. 1b, the data can be considered as horizontally partitioned in federated learning, when different data are generated on

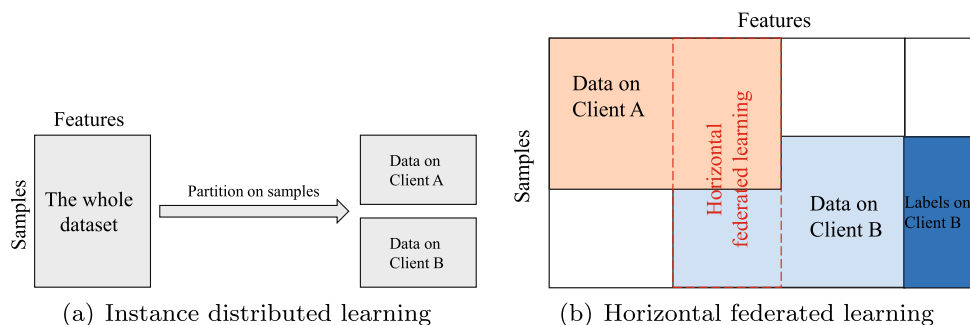
different clients that have the same attributes (features). For instance, two hospitals in different regions may have different patients, although they performed the same tests for each patients and collected the same personal information such as the name, age, gender and address.

There are three main differences between instance distributed learning and horizontal federated learning. First, data are typically independent and identically distributed (IID) in distributed learning but may be non-IID in horizontal federated learning. As mentioned before, distributed learning is mainly designed for reducing the training time; therefore, designers can manually allocate every subsets of the client data to be IID to enhance the convergence. However, in horizontal federated learning, the central server has no access to any raw data, which are usually non-IID on different clients. Second, horizontal federated learning always contains massive connected clients, but instance distributed learning often does not have a very large number of workers, because too many workers will worsen the performance of distributed training, when the total amount of data is fixed [5]. Finally, global model update mechanisms are slightly different. In instance distributed learning, such as multi-GPU training (GPUs are always embedded inside a server, communication effects can thus be ignored), a deep neural network tends to synchronously update the global model once the local gradients of the mini-batch data are calculated, to ensure the correct distributed model learning direction. This global model updating approach is intrinsically not suited for horizontal federated learning because frequent upload and download of data are not desirable due to the constraints on the communication costs.

Typical horizontal federated learning (Fig. 2) algorithms, such as the FedAvg proposed in [3], consist of the following main steps.

1. Initialize the global model parameters on the server and download the global model to every participating (connected) clients.
2. Every connected clients learn the downloaded global model on its own data for several training epochs. Once completed, the updated model parameters or gradients

Fig. 1 Instance distributed learning (a) and horizontal federated learning (b)






Fig. 2 Flowchart of federated learning. θ is the global model parameters, n_k is the data size of client k , K is the total number of clients and t is the communication round in federated learning. We just initialize global model parameters randomly at the beginning of the communication round and use updated model parameters afterwards

(gradients here means the difference between the downloaded model and updated model) would be sent to the server. Note that the clients may have different amounts of training data and unbalanced computational resources. As a result, the server is not able to receive the uploads from different clients at the same time.

3. The server aggregates the received uploads (synchronously or asynchronously) to update the global model.
4. Repeat the above two steps until convergence.

From the above steps, we can find that the central server can only receive model weights or gradients of the participating clients and has no access to any local raw data. Therefore, users’ privacy is immensely protected in horizontal federated learning.

The above steps of the horizontal federated learning algorithm appear to be very similar to those in synchronous distributed systems. However, these two schemes have different learning environments and purposes. Horizontal federated learning is performed in a more complex environment since the connected edge devices like mobile phones may become frequently offline and the global learning performance cannot be guaranteed. In contrast to federated learning, distributed learning systems are often designed and run in a more stable environment to ensure a good learning performance. Apart from that, the purpose of federated learning is to protect local user’s private data, while distributed learning is mainly designed to accelerate the learning speed.

Horizontal federated learning has three additional main challenges compared to the traditional centralized learning:

- (1) it must reduce the communication resources as much as possible,
- (2) it needs to improve the convergence speed, and
- (3) it must make sure that no private information is leaked in passing the model information.

Much research work has focused on reducing communication costs, such as client updates sub-sampling [6–8] and model quantization [9,10]. More recently, Chen et al. [11] propose a layer-wise asynchronous update algorithm to reduce the communication costs by decreasing the update frequency of the deep layers in the neural network. In addition, Zhu et al. [12] use a multi-objective evolutionary algorithm (MOEA) to simultaneously enhance the model performance and communication efficiency. Learning a good model in horizontal federated learning is not an easy task since the training data on different clients are usually non-IID, leading to possible model divergence. To solve this issue, Zhao et al. [13] empirically explore the effect of non-IID data and provide a statistical analysis of divergence. Li et al. [14] propose a FedProx algorithm to alleviate negative impacts of the system’s heterogeneity by injecting a proximal term into the original loss on each client. Apart from it, an attentive aggregation method [15] is used to minimize the weighted distance between the server model and client models on non-IID datasets.

The central server is often regarded as honest but curious (follow federated learning protocol but try to infer client data information) in horizontal federated learning, and the revealed gradients of each client may potentially leak the data information [6]. For this reason, Phong et al. [16] mathematically prove that model gradients (especially the first hidden weights) are proportional to the original data and adopt additive homomorphic encryption [17] to encrypt and protect model gradients. In their method, the secret key is kept confidential to the server but known to all participating clients and the central server can easily get the plain model gradients as long as one of connected clients uploads its secret key. In order to mitigate this issue, secure multiparty computation (SMC) [18,19] is proposed to partition an intact secret key into several key shards and each client can just hold one shard. As a result, the server must get at least t shards (t is the threshold value) for decryption. Consequently, privacy preserving is significantly improved.

However, homomorphic encryption will increase the computation load, and SMC consumes much more communication resources, since encrypted model weights need to be downloaded and uploaded between the server and at least t clients for partial decryption. Therefore, a more light-weight privacy-preserving technique called differential privacy [20] can also be used in horizontal federated learning. Such as the methods used in [6,21], a Gaussian or Laplacian noise is added to the gradients of each client before sending them to the central server. Note, however, that the learning process may be interrupted if the accountant [22] exceeds a pre-defined threshold value. Most recently, Truex et al. suggest a

hybrid approach that combines differential privacy together with homomorphic encryption.

Vertical federated learning

In contrast to horizontal federated learning, vertical federated learning is applicable to the situations where the datasets share the same sample space but have different feature space, as shown by part of surrounded by the dashed lines in Fig. 3b. For example, two different financial agents may have the same customers but provide different services. Different from horizontal federated learning, vertical federated learning is similar to feature distributed learning [4] to some extent which 'vertically' partitions the training data, as shown in Fig. 3a upon the feature space. Moreover, the central server is often called a coordinator [23] in feature distributed learning or vertical federated learning since its main task is to calculate the total loss rather than aggregating the uploaded weights.

Vertical federated learning is first introduced in [23], in which the overall framework contains one trusted coordinator and two parties, where each party represents one client. The coordinator computes the training loss and generates encryption key pairs. Homomorphic encryption is adopted for the privacy-preserving purpose and the effect of entity resolution is also discussed. More recently, a two-party architecture [24,25] is proposed by removing the trusted coordinator which greatly reduces the complexity of the system. A typical two-party framework of vertical federated learning using a simple logistic regression model includes the following steps:

1. Assume Party A contains the data labels. Party A creates a homomorphic encryption key pair and sends the public key to Party B. Both parties initialize their local model parameters according to their feature dimensions of local training data.
2. Both parties compute their local inner products of data and the model. Then Party B sends its results to Party A.
3. Party A sums two inner products and calculates the loss function by data labels. The loss is encrypted with a public

key and is sent to Party B. The model gradients of Party A are also calculated.

4. Party B calculates the encrypted model gradients from the received loss and encrypt. In addition, a random number is encrypted and added to the encrypted gradients. The summation should be sent to Party A for decryption.
5. Party A uses a secret key to decrypt the summation value and sends it to Party B.
6. Update both model parameters of the two parties.
7. Repeat Step 2 to Step 6 until convergence.

In Step 3, the training loss is encrypted before being sent to Party B, because it contains the information of the data labels in Party A which cannot be revealed to Party B. As a result, Party B needs to calculate its local model gradients on the encrypted loss and a Taylor approximation is commonly used [23,25] for simplifying this computation.

From the above discussions, we can see that vertical federated learning is dramatically different from horizontal federated learning. The central server in horizontal federated learning is used for model aggregation, while in vertical federated learning the server plays the role of calculating the loss or collecting features. In addition, the server can be removed in vertical federated learning, e.g., summing the training loss within one of participating parties (clients). Apart from the above, we often assume that not all parties contain the data labels in vertical federated learning, e.g., only Client B contains data labels in Fig. 3b and those parties with no data labels are not able to update their models locally. Therefore, we call the server 'coordinator' that coordinates the feature predictions from all parties for calculating the training loss in vertical federated learning.

Most studies of vertical federated learning only support two parties (with or without a central coordinator) using a simple binary logistic regression model. Feng et al. [26] adopt the idea of multi-view learning to extend the previous scheme into a multi-participant multi-class vertical federated learning framework. Besides, Liu et al. introduce a federated stochastic block coordinate descent algorithm, where all participating parties update their local models for multiple times to reduce the total number of communication rounds. In addi-

Fig. 3 Feature distributed learning (a) and vertical federated learning (b)

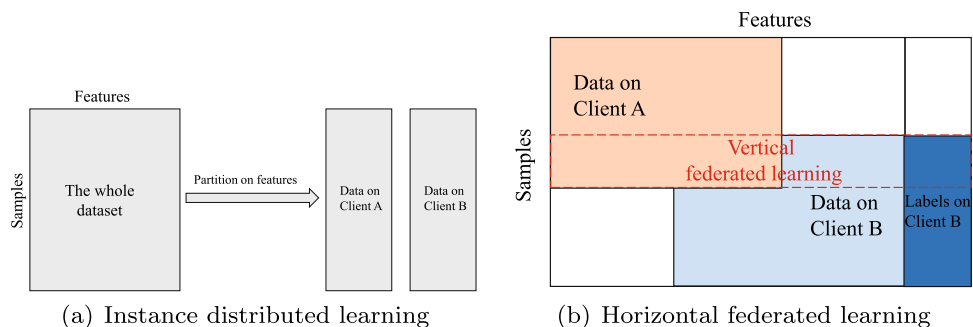
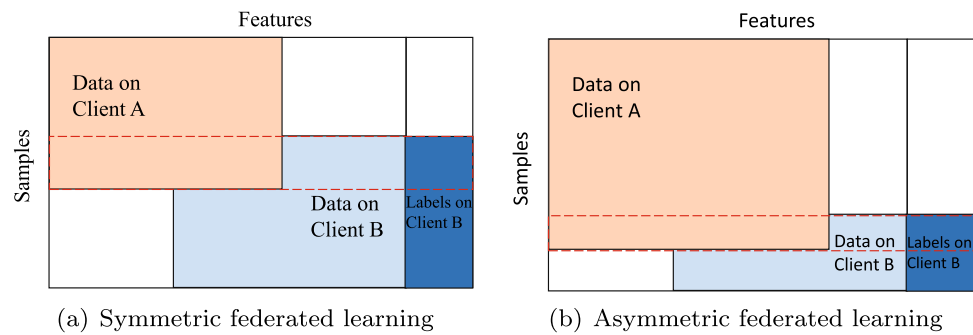


Fig. 4 **a** Symmetric federated learning, and **b** asymmetric federated learning



tion, Chen et al. propose an asynchronous vertical federated learning method and differential privacy is also used for privacy protection.

Hybrid federated learning

Hybrid federated learning is more realistic in the real world and it assumes that datasets on different clients not only have different sample spaces but also different feature spaces. Therefore, in this scenario, different parties need to share the data identity (ID) information to find the intersection part for distributed training, which is a threat to local clients' privacy. Since participants in hybrid federated learning are often asymmetric [27], for instance, some participants are small companies always requiring to protect their ID information, while some participants are large companies that have no concern about the ID privacy. Symmetric federated learning and asymmetric federated learning are illustrated in Fig. 4.

Secure ID alignment protocol is significant for hybrid federated learning, such as the commonly used Private Set Intersection (PSI) protocol. In standard PSI, all participants want to collaboratively find the intersection (the part indicated by the dashed lines in Fig. 4) and keep unintersected parts private. The PSI protocols can be implemented by a classical public-key cryptosystem [28,29] or other similar techniques.

Federated model training is similar to vertical federated learning; however, for asymmetric federated learning, Genuine with Dummy (GWD) approach [27] is used to ensure the correctness of computation results.

Neural architecture search

Since the quality of deep neural networks (DNNs) heavily depends on their architecture, increasing research efforts have been committed to design of novel structures in the deep learning community. However, manually designing deep neural networks requires considerable expertise in the field of deep learning and the investigated problem, which is unrealistic for many interested users. Not until recently has

automated machine learning (Auto ML), in particular neural architecture search (NAS) become very popular to allow interested users without adequate domain knowledge to profit from the success of deep neural networks. The framework of NAS methods involves three dimensions [1], namely search space, search strategies, and performance estimation strategies.

The search space is a collection of network architectures, which has a major influence on the performance of the generated networks and search efficiency. The search strategy defines the method used to automatically design the optimal network architecture. To be specific, these search strategies can be divided into at least three categories: reinforcement learning (RL), evolutionary algorithms (EAs), and gradient-based (GD) methods. In addition, a few additional methods, such as random search [30,31], Bayesian optimization [32,33] and multinomial distribution learning [34], fall outside of the above categories. The search strategy aims at finding architectures that can obtain high performance on the test dataset. To guide searches effectively, these strategies utilize a performance estimation strategy to evaluate the quality of candidate architectures. Early work uses a simple way of performance estimation, for example, by iteratively training a candidate architecture on the training dataset with the stochastic gradient descent (SGD) [35] and evaluating its performance on the validation data [36–41]. Such an evaluation strategy typically results in a prohibitively high computational cost. For example, to design a good performance of neural network, the automatic evolving convolutional neural network (AE-CNN) [40] algorithm consumes 27 GPU-days and the neural architecture search approach [36] consumes 22400 GPU-days on the CIFAR10 dataset. Because inefficient search strategies require a large number of GPUs, many NAS methods cannot be implemented given limited computational resources. To address these challenges, much recent work dedicates to developing effective methods which can reduce the computational costs of performance evaluation, e.g., surrogate-assisted evolutionary algorithms (SAEAs) [33,42,43], information reuse [44,45], one-shot neural architecture search [46–50], among many others.

NAS based on reinforcement learning

Early work on NAS depends on RL to search for high-performance neural architectures [36–38]. The design of a network model is considered as an agent's action, which specifies the architecture of the network (i.e., a child model). The network is then trained and its performance on the validation data is returned as the agent's reward.

A policy gradient method has attempted to approximate some nondifferentiable reward function to train a model which needs parameter gradients. Zoph et al. [36,38] first adopt this algorithm in NAS to train a recurrent neural network (RNN) model that generates architectures. As is shown in Fig. 6, the controller as a navigating tool to find more promising architectures in the search space. The original method in [36] uses a macro-search space that generates the entire network at once. As is shown in Fig. 5, the whole architecture consists of n sequential layers where the dashed lines indicate skip connections. Hence, the macro-search space aims to design the entire CNN architecture in terms of the number of hidden layers n , operations types (e.g. convolution), network hyper parameters (e.g., the convolutional kernel size), and the link methods (e.g. skip connections). However, this method is expensive when the data set is large. To reduce the computational cost, Zoph et al. [38] propose a more structured search space, called micro-search space. The micro-search space only covers repeated smaller modules, called normal cell and reduction cell, and then connects them together to form an entire network. As shown in Fig. 5, these cells are built in complex multi-branch operations (e.g. convolution). Each cell structure contains two inputs

$h[i - 1]$ and $h[i - 2]$ coming from two previous layers. Hence, the micro-search space aims to design structures of these two types of cells. In addition, the cell structures should have a good capability of generalizing to other related tasks. For example, the proposed method searches for optimal cell structures on the CIFAR10 data set and transfers them to the ImageNet data set by stacking together multiple copies of this cell. After that, the NASNet [38] method is extended to a multi-objective optimization variant to simultaneously optimize the classification performance and computational cost using different scalarization parameters.

Q-learning [51], as a class of popular RL methods, has been widely used for NAS. Baker et al. [52] employ an ϵ -greedy Q-learning strategy to train a policy that sequentially chooses a type of layers (e.g. convolutional layer, pooling layer, and fully connected layer) and their corresponding hyperparameters. Zhong et al. [53] extend this method to a block-wise network generation approach, which designs blocks with the same Q-learning paradigm. After that, the optimal blocks are repeated and stacked to construct the entire network architecture. To accelerate the search process, a distributed asynchronous strategy and an early-stop approach are adopted.

Parameter sharing introduced in efficient NAS (ENAS) [46] is a promising approach for speeding up the search process for RL-based NAS methods, which treats architectures as different sub-graphs (sub-net) of a larger graph (super-net) and forces all sub-graphs to share a common set of weights that have edges of this larger graph in common. Pasunuru et al. [54] propose a multi-task architecture search (MAS) approach based on ENAS [46] for finding a cell structure that

Fig. 5 Illustration of the macro- and micro-search spaces

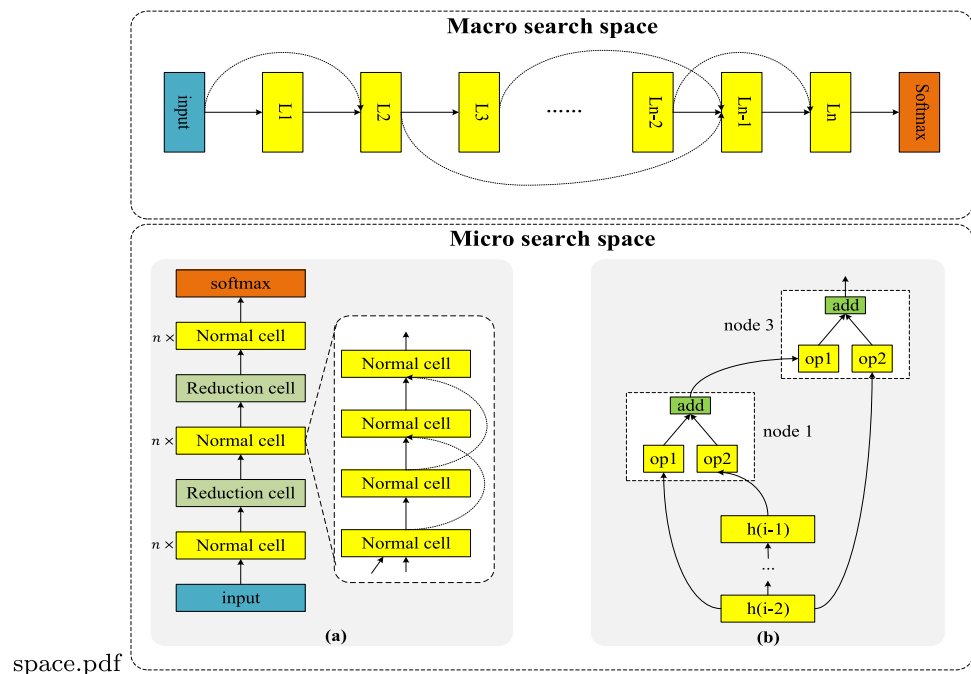
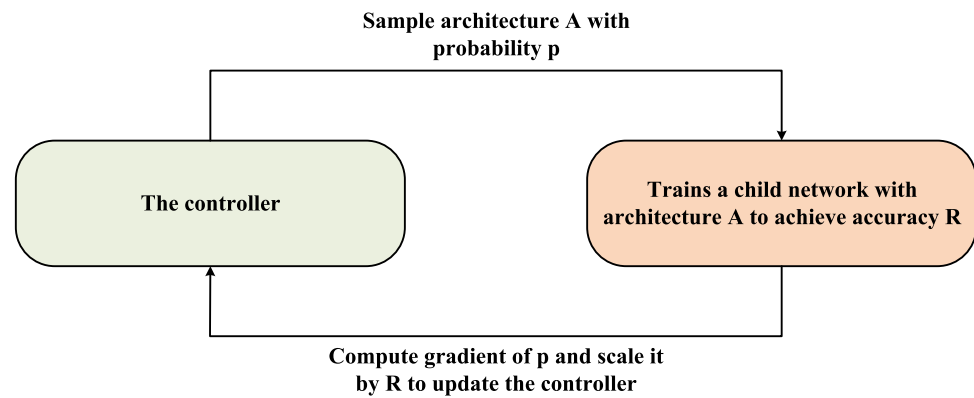


Fig. 6 An overview of RL-based NAS method



performs well across multiple tasks. Hence, the cell structure generated by NAS can transfer to a new task. Bender et al. [55] propose a thorough comparison between random search NAS methods and ENAS [46] on a larger search spaces for image detection and classification tasks, respectively. In addition, a new reward function is suggested, which can effectively improve the quality of the generated networks and reduce the difficulty for manual hyperparameter tuning. Liu et al. [56] present a novel knowledge distillation [57] approach to NAS, called architecture-aware knowledge distillation (AKD), which finds student models (compressed teacher models) that are best for distilling the given teacher model. The authors employ a RL-based NAS method with a KD-guided reward function to search for the best student model based on a given teacher model.

NAS based on EAs

EAs are a class of population-based, gradient-free heuristic search paradigms, which have been widely used in solving various complex optimization problems [12,58–60]. Historically, EAs have already been used for simultaneous optimization of the topology, weights of the connections and hyperparameters of artificial neural networks (ANNs) [61,62,62–64]. The neuroevolution with augmenting typologies (NEAT) algorithm [65] is one of the popular early methods that have shown powerful performance. However, the traditional approaches are not well suited for optimizing DNNs due to the complex network architectures and large quantities of connection weights. EA-based NAS approaches to optimizing deep network architectures have started gaining momentum again recently [66,67], mainly because they can simultaneously explore multiple areas of the search space and their relative insensitiveness to a local minimum [68,69]. Figure 7 shows a generic framework of EA-based NAS algorithms. Broadly speaking, the whole process of an EA-based NAS algorithm follows the procedure of an EA containing at least four-steps: population initialization, offspring generation, fitness evaluation, and environmental selection.

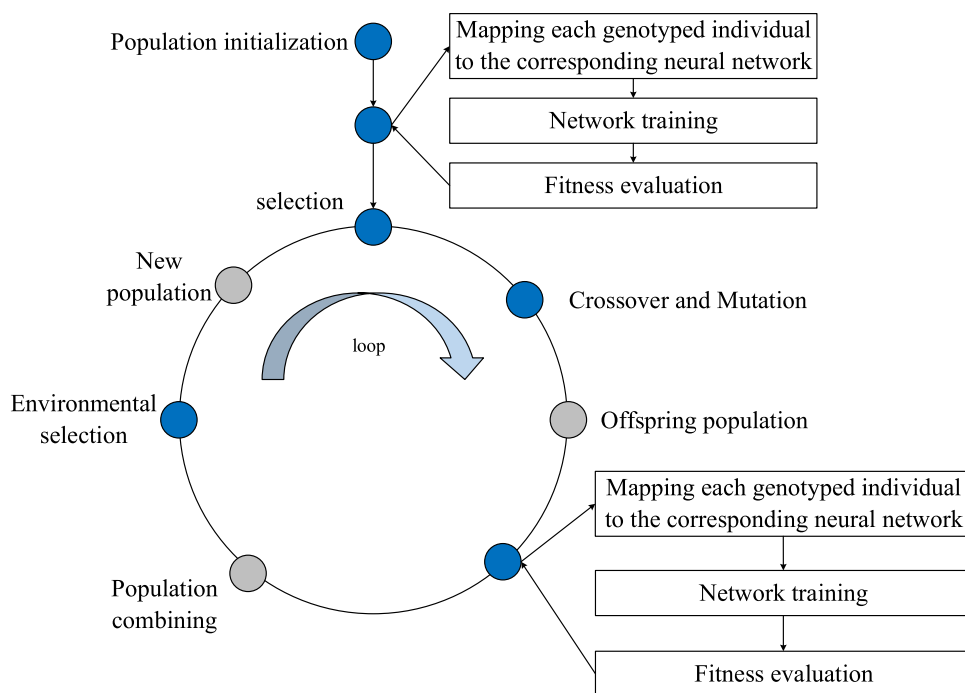
Generally, each neural network in the search space is encoded as a chromosome, and crossover and mutation operations of the chromosomes are performed in the exploration. Then each chromosomes is transformed into a corresponding neural network, and iteratively trained on the training dataset. The trained network is evaluated on the validation dataset to get their fitness value.

Xie et al. proposed a genetic CNN [39] method that is one of the early studies using an EA for optimizing convolutional neural networks (CNNs). The genetic CNN algorithm searches over the entire architecture space and employs a fixed-length binary string to represent the connection between a number of ordered nodes (e.g. a convolutional operation). Although this early algorithm has some limitations, including a limited number of nodes as well as limited sizes and operations of convolutional filters, the generated structures have not only achieved competitive results on the CIFAR and SVHN datasets, but also shown excellent transferability to the ImageNet dataset [70].

Miikkulainen et al. [71] propose a coevolution DeepNEAT (CoDeepNEAT) method by extending the NEAT algorithm [65] to DNNs. In CoDeepNEAT, each neural network is assembled by modules and blueprints. A coevolutionary method is adopted that evolves two populations of modules and blueprints separately, in which each module chromosome represents a small DNN. The blueprints chromosome represents a graph where each node contains a pointer to a particular module species. The assembled networks are trained and evaluated in an ordinary way of NAS. The fitness of the network is the average fitness of the entire candidate models containing the blueprints or modules. In addition, Liang et al. found that the CoDeepNEAT also achieves promising performance in the Omniglot multi-task learning domain [72].

In fact, the length of a chromosome usually represents the depth of the corresponding neural network and a fixed encoding scheme may limit the performance of the optimized network. To address this issue, some recent NAS algorithms have attempted to use a variable-length encoding scheme.

Fig. 7 A generic EA-based NAS framework



Real et al. [73] propose a large-scale evolutionary NAS method, which utilizes a variable-length encoding method in which the network architectures can adaptively change their depths. Sun et al. [40] propose an AE-CNN algorithm that can fully automatically design CNN architectures, without requiring any pre-processing or post-processing. Inspired by the ResNet [74] and DenseNet [75], AE-CNN's search space is defined by some predetermined building blocks, including ResNet block and DenseNet block, max pooling layer and mean pooling layer. Then, the authors design an EA-based NAS framework, including the variable-length encoding and a novel crossover and mutation operators based on the variable-length encoding, as the search strategy to search the optimal depth of the CNN. Given the nature of the variable-length encoding strategy, the algorithm employs a repair mechanism that avoids to produce invalid CNNs. Inspired by directed acyclic graph (DAG), William et al. [76] introduce a DAG-based encoding strategy, which can represent CNNs of an arbitrary connection structure and an unlimited depth.

Suganuma et al. [77] propose a CGP-CNN algorithm to design CNN architectures using genetic programming. The search space of CGP-CNN is represented by a DAG, where the nodes represent either convolutional blocks or concatenation operations. Then CGP-CNN uses the Cartesian genetic programming (CGP) [78,79] encoding scheme to represent network architectures and their connectivity. This encoding scheme can represent variable-length network architectures and skip connections.

Most EA-based NAS methods aim at finding better topologies for DNNs while leaving the learning of weights to SGD. It is known that the SGD optimizer heavily relies on the initial values of the weights. To alleviate this problem, Sun et al. [80] propose an EA-based NAS method, named Evolving Deep CNNs (EvoCNN), to automatically design CNN architectures and corresponding connection weight initialization values without manual intervention. To reduce the search space, two statistical measures, including the mean and standard deviation of the connection weights, are encoded in the chromosome to represent tremendous numbers of the connection weights. In addition, the incomplete training scheme is employed to accelerate the fitness evaluation. According to the Occam's razor theory [81], the number of connection weights is also considered as an indicator to scale the quality of candidate networks.

Sun et al. [41] use a genetic algorithm (GA) to design CNN architectures (CNN-GA). In CNN-GA, the standard convolutional layer is replaced by a novel building block, called the skip layer. The skip layer consists of two convolutional layers and one skip connection. Hence, the genotype encodes information of the skip layers and pooling layers. The fully connected layers are discarded, mainly because they easily result in the overfitting [82].

Rather than generating the entire CNNs, the micro-search space [46] has also been successfully employed by many recent EA-based NAS algorithms [83–87]. Real et al. [85] propose an extension of the large-scale evolution [73], called AmoebaNet, which has achieved better results on ImageNet compared with hand-designed methods for the first time.

Since EAs are a class of population-based search methods, the main computational bottleneck of EA-based NAS approaches lies in evaluating the fitness of the individuals by invoking the lower-level weight optimization. One such evaluation typically takes several hours to days if the network is large and if the training dataset is huge. For instance, on the CIFAR10 datasets, the AE-CNN [40] consumed 27 GPU days, CNN-GA [41] consumed 35 GPU days, and the large-scale evolutionary algorithm [73] consumed 2750 GPU days, AmoebaNet [85] consumed 3150 GPU days. This seriously limits the practical usability of most evolutionary NAS methods under a constrained search budget.

Therefore, various techniques have been suggested to accelerate the fitness evaluation, such as information reuse [44,47] and SAEAs [88]. SAEAs have been popular to solve computationally expensive optimization problems, which use cheap classification and regression models, e.g., radial basis function networks (RBFNs) [89,90] and Gaussian process (GP) models [91,92], to replace the expensive fitness evaluation [93]. Generally, the candidate networks are trained from a few number of expensive fitness evaluations, and then the trained networks are used to build a fitness predictors to reduce the cost of fitness evaluations. In the area of evolutionary NAS, Swersky et al. [33] adopt Bayesian optimization [94] to speed up evolutionary optimization, which is called freeze–thaw Bayesian optimization. Unfortunately, this algorithm is based on Markov chain Monte Carlo sampling and also suffers from high computational complexity. Sun et al. proposed a performance predictor termed E2EPP, which is based on a class of SAEAs method [43] meant for offline data-driven evolutionary optimization of expensive engineering problems. Specifically, E2EPP builds a surrogate that can predict the quality of a candidate CNN, thereby avoiding the training of a large number of neural networks during the search process. Compared with AE-CNN, a variant of AE-CNN assisted by E2EPP (called AE-CNN+E2EPP) can achieve a 2.3x and 2.1x reduction in GPU days on CIFAR100 and CIFAR10, respectively. Lu et al. [95] adopt two surrogates to address the bi-level NAS problem. On the one hand, a fine-tuning method is adopted at the weight level to improve the efficiency of SGD training. On the other hand, an online learning method is used to improve the sample efficiency for the search space at the neural architecture level.

Knowledge inheritance [44,47] is another promising approach to accelerate fitness evaluations. Zhang et al. [44] propose an EA based on asexual reproduction to find better typologies for deep CNNs and knowledge inheritance to reduce the computation cost. Once the topology of an offspring individual is generated by its parent, the weights of offspring networks are directly copied from its parents. For edges that do not appear in its parent network, the weights are randomly initialized.

To reduce the computational burden for fitness evaluations, another widely adopted approaches are to train and evaluate individuals using proxy metrics [85–87]. The performance of the proxy models is used as the surrogate measurements to guide the evolutionary search. Such proxy metrics include reducing the width (the number of channels) and the depth (number of layers) for the intended network architecture to create a small-scale network, shortening the training time, reducing the resolution of input images, and training on a subset of the full training dataset. However, these simple proxy model constructing methods may result in a low correlation in prediction mainly because they may introduce biases in fitness estimation. Zhou et al. [86] have conducted extensive experiments on different combinations of proxy metrics to investigate their behaviors in maintaining the rank consistency in NAS, based on which a reliable hierarchical proxy strategy is proposed to accomplish economical neural architecture search (EcoNAS). The hierarchical proxy strategy aims at discarding less promising candidate individuals earlier with a fast proxy and estimates more promising individuals with a more expensive proxy. Hence, the EcoNAS method is able to achieve a 400× reduced search time in comparison to AmoebaNet [85] without sacrificing the performance. Lu et al. [87] empirically establish the trade-off between the correlation of proxy performance to true performance and the speed-up in estimation.

Evolutionary multi-objective NAS methods considering multiple conflicting objectives have been reported. One of the earliest evolutionary multi-objective methods to design CNNs is NEMO [96], which simultaneously optimizes classification performance and inference time of a network based on NSGA-II [97]. Inspired by NEMO, Lu et al. [84] consider classification error and computational complexity as the two objectives. In addition, they empirically test multiple computational complexity metrics to measure the inference time of a network containing the number of active layers, the number of activating connections between layers, and the number of floating-point operations (FLOPs). And then, the FLOPs are used as a second conflicting objective for optimization. Moreover, a Bayesian network (BN) is adopted to learn the knowledge about promising architectures present in the search history and then guide the future exploitation in generating new architectures. Subsequently, Lu et al. suggest NSGANet-v2 [87], an extension of NSGANet [84], where a comprehensive search space including more layer operations and one more option that controls the width of the model is introduced. Dong et al. [98] present a DPP-Net on the basis of [99] that optimizes both GPU memory usage and the model performance. Elsken et al. [100] proposed the LEMON-ADE method, which formulates the NAS as a bi-objective optimization problem that maximizes the performance and minimizes the required computational resources. Inspired by [101], LEMONADE reduces computational cost through

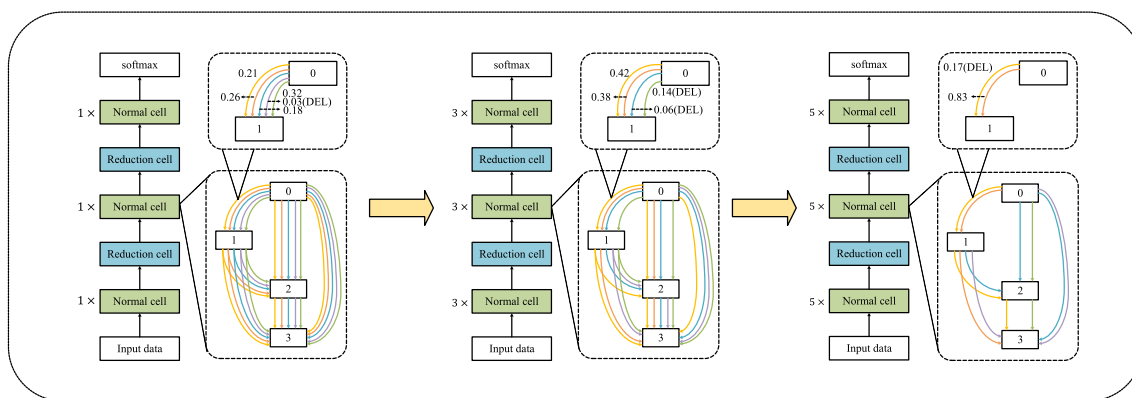


Fig. 8 A generic pipeline GD-based NAS method

a custom-designed approximate network morphisms, which makes offspring individuals to share weights with their fore-runners, avoiding training new networks from scratch. Note that evolutionary multi-objective structure optimization of shallow networks can be traced back to a decade ago [102]. Lu et al. [103] proposed a method that integrates multi-objective evolutionary algorithm and online transfer learning for designing task-dependent network architectures trading-off model performance and computational complexity. The authors first train a task-specific one-shot model covering the search space once, then specialized sub-models can be sampled from the search space without additional training. And then they adopt an online regressor as a surrogate model to predict the performance of sub-models in the one-shot model.

NAS based on GD

Compared with the above gradient-free optimization methods, the GD-based methods (Fig. 8) have become increasingly popular recently, mainly because their search speed is much faster than RL-based and EA-based methods. Early GD-based methods [104–107] implement this idea for optimizing layer hyperparameters or connectivity patterns, respectively. Lorraine et al. [108] introduce an algorithm for inexpensive GD-based hyperparameter optimization. Liu et al. [109] employ GD in the DARTS algorithm, which optimizes both the network weights and the architecture. The authors use relaxation tricks to make a weighted sum of candidate operations differentiable, and then apply the gradient descent method to directly train the weights. Inspired by DARTS [109], Dong et al. [110] introduce gradient-based search using the differentiable architecture sampler (GDAS) method. The authors develop a differentiable architecture sampler which samples individual architectures in a differentiable way to accelerate the architecture search procedure. Stochastic NAS (SNAS) [111] optimizes a probability distribution of the connections between different candidate operations. Li et al. [112] observe that models with a higher

performance during the search phase may perform worse in the evaluation. Hence, they divided the search process into sub-problems and proposed sequential greedy architecture search (SGAS) based on DARTS, which chooses and prunes candidate operations (e.g. convolutional layers) greedily. The authors apply SGAS for CNNs and graph convolutional networks (GCNs) and have achieved competitive performance. Xu et al. [113] present Partially-Connected DARTS (PC-DARTS), which samples a small part of super-network to reduce the redundancy in exploring the network space. Compared with DARTS, PC-DARTS not only enjoys both faster speed and higher training stability but also a highly competitive learning performance. Gao et al. [114] propose the first GD-based NAS method in generative adversarial networks (GANs), called adversarial neural architecture search (AdversarialNAS), which can search the architectures of generator and discriminator simultaneously in a differentiable manner.

One bottleneck of the above GD-based NAS methods (e.g. DARTS) is that it requires excessive GPU memory during search in that all candidate network layers must be explicitly instantiated in the GPU memory. As a result, the size of the search space is constrained. To address this issue, Wan et al. [115] propose DMaskingNAS, a memory and computationally efficient DARTS variant. DMaskingNAS employs a masking mechanism for feature map reuse. Hence, although the search space of DMaskingNAS is expanded up to $10^{14} \times$ over conventional DARTS, memory and computational costs stay nearly constant.

Another way to address the above problem is to utilize proxy tasks, e.g., learning with only a small number of building blocks or training for a small number of epochs [109, 111]. However, these approaches cannot guarantee to be optimal on the target task due to the restricted block diversity [107]. Cai et al. [107] proposed ProxylessNAS method, which directly designs the networks based on the target task and hardware instead of with proxy. Meanwhile, the authors used path binarization to reduced the computational cost (GPU-hours and

GPU memory) of NAS to the same level of normal training. Hence, ProxylessNAS algorithm can generate network architectures on the ImageNet dataset without any proxy.

Most recently GD-based NAS methods are formulated as bilevel optimization problems. However, He et al. [116] observe that bilevel optimization in the current methods is solved based on a heuristic. For instance, solution of the problem needs to get an approximation of the second-order methods [109,110]. He et al. [116] demonstrate that the approximation has a superposition influence mainly because it is based on a one-step approximation of the network weights. As a result, gradient errors may cause the algorithm to fail to converge to a (locally) optimal solution. Hence, the authors propose mixed-level reformulation NAS (MiLeNAS) that uses a first-order method on the mixed-level formulation. Experimental results show that MiLeNAS has achieved higher classification accuracies than those achieved by the original bilevel optimization methods.

Federated neural architecture search

Federated NAS aims to optimize the architecture of neural network models in the federated learning environment. As discussed in section “Federated Learning”, distributed model training is intrinsically more difficult than centralized training, and it becomes even more challenging for NAS problems. In this section, we would like to introduce the current research on federated NAS and discuss them from two perspectives: online and offline optimization, and single- and multi-objective optimization. It should be noticed that research on federated NAS work is presently limited to horizontal federated learning and federated NAS in vertical federated learning has not been reported so far.

Offline and online federated neural architecture search

Most NAS methods include two steps, i.e., search the architecture of the neural network model, and training the weights of the found neural network model afterwards. And most importantly, only the final performance matters. We define these approaches as offline NAS, because the search and training steps are typically separate and only an optimized network will be used. By contrast, online NAS requires that the architecture optimization and weight training be done at the same time, and some of the models must be used during the search process. As a result, the performance of the models during the optimization must be acceptable.

This concept can be easily extended to federated learning. In other words, federated NAS systems in which neural architecture search and weight training of the global model must be performed simultaneously are called online or real-time

federated NAS, whilst federated NAS in which neural architecture search can be conducted at first and then the weights of the found models are trained are offline. Similarly, online federated NAS requires that the neural network models can be used during the optimization process.

For example, the method proposed in [12] is a typical offline federated NAS framework using a multi-objective evolutionary algorithm. An offline evolutionary federated NAS algorithm can be summarized as follows:

1. Initialize parents with a population size N and each individual represents one architecture of the neural network. Construct and train N neural network models in federated learning with all participating clients to achieve the fitness values (e.g., validation accuracy) of the parents.
2. Generate N offspring individuals by applying genetic operators on the parents. Construct and train all the generated offspring models for fitness evaluations in federated learning.
3. Combine the parent and offspring populations into one population and perform environmental selection. Select the best N individual from the combined population as the new parents.
4. Repeat the above two steps until the evolutionary algorithm converges.
5. Train the weights of the optimized neural network models in federated learning.

It can be seen that all participating clients are used for federated model training, i.e., at each generation, all participating clients must train each of the N individuals for certain rounds for fitness evaluations, which significantly increases both computation and communication costs. Client sampling can be used to alleviate this issue, in which only subsets of participating clients contribute to one individual’s model training. For example in [117], all the connected clients are divided into different groups and each sampled model use one group of clients for local training. The overall process of this approach is summarized as follows:

1. Initialize the global model and a list of resource budgets in the server.
2. Generate a list of simplified global models by model pruning [118] based on the current global model. And then these global models are distributed to different group clients.
3. For each communication round, every group of clients train their allocated group models for a number of predefined epochs and calculate the test accuracies on the validation datasets. Then both local test accuracies and validation data sizes are uploaded to the server. The server aggregates the uploaded local models and calculates a weighted accuracy for each group model. Remove $\alpha\%$

of global models with the worst test accuracies (removed global models that will not be trained and updated from the next communication round). The remaining groups of clients upload their calculated model gradients to the server for aggregation.

4. Repeat the above step for a number of pre-defined communication rounds.
5. Replace and store the global model with the first model in the global model list.
6. Repeat the above four steps until convergence.
7. Perform federated training on any stored global model.

Although the above procedure uses different architecture generation methods (model pruning) and search space compared to the evolutionary approach, it is clearly a population based offline federated NAS framework (weight training and architecture search are separate). In addition to client sampling, the authors also remove subsets of global models to further reduce the communication costs. However, the test accuracies of each global model in the list is calculated before model aggregation, which sometimes cannot represent the real test accuracies, especially for the cases when the clients' data are particularly non-IID.

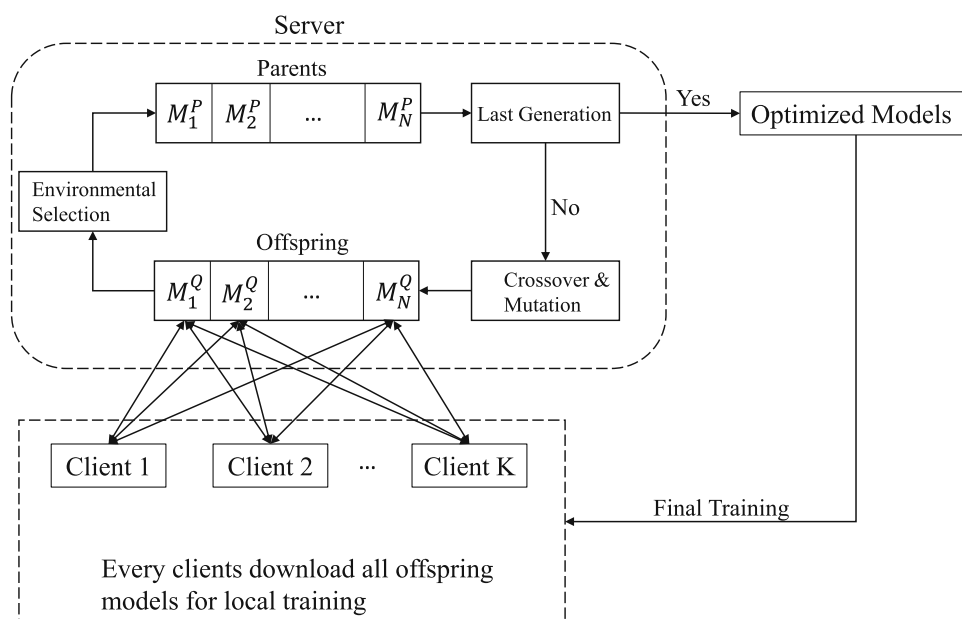
The overall framework of Offline federated NAS is shown in Fig. 9 and it has two main difficulties: (1) the number of communication rounds for federated model training of each individual is hard to determine. Setting a small number of communication rounds may make the individual's model under-trained and bias the fitness evaluations. On the other hand, setting a very large number of communication rounds consumes too many communication resources. (2) Training the candidate neural network models consume additional

communication resources, which should be avoided in federated learning. For the above reasons, online federated NAS frameworks need to be developed to solve the above issues.

Online federated NAS trains the model and does the architecture optimization simultaneously (shown in Fig. 10). To the best of our knowledge, there are currently two approaches to online federated NAS. One is gradient-based method proposed in [119], and the other is an EA-based method proposed in [120]. The gradient-based method adopts the idea of DARTS [109], which is implemented in the federated environment. The global model here is called supernet which consists of repeated directed acyclic graphs (DAGs) and each DAG contains all candidate operations. And relaxation tricks [121] are used to make a weighted sum of the candidate operations differentiable so that the architecture parameters can be directly updated by the gradient descent algorithm. A brief description of this method is given below.

1. The server initializes the supernet and its architecture parameters.
2. All connected clients download the supernet and its architecture parameters from the server.
3. Each client trains and updates the supernet with fixed architecture parameters on mini-batch training data at first. Then update the architecture parameters with fixed model parameters on mini-batch validation data. These two procedures are performed alternately within one local training epoch.
4. After local training for several epochs, all participating clients upload both model and architecture parameters to the server. The server performs weighted averaging upon the supernet and architect parameters.

Fig. 9 Overall framework of offline federated NAS






Fig. 10 Overall framework of online federated NAS. Client sampling are often used to ensure all offspring models are evaluated within one communication round

- Repeat from step 2 to step 4 until convergence.

Unlike aforementioned two offline federated NAS framework, this scheme is not population based since all candidate operations are jointly optimized. In addition, architecture search and weight training of the supernet model are conducted alternately during the period of federated training. Therefore, no additional communication resources are required for training the candidate models. However, jointly optimizing the supernet on local clients requires much more computation and memory resources, which is not well suited for the edge devices like mobile phones.

To reduce the memory usage of local devices, a more light-weighted real-time evolutionary NAS framework (RT-FedEvoNAS) is proposed in [120]. Different from the previous gradient-based approach, RT-FedEvoNAS adopts model sampling technique [47,122] in federated learning, where only one path of repeated cells in the global model is sampled and downloaded to local clients. As a result, both local computation and uploading costs are significantly reduced. The overall process is described as follows:

- Initialize the supernet in the server. Generate the parent population containing N individuals, each representing a one-path subnet sampled from the supernet using a choice key. Do client sampling to allocate L clients evenly into N groups.
- Download the subnet of each parent individual to each group of clients for training. Once the training is com-

pleted, upload the L local subnets to the server for aggregation to update the supernet model.

- Generate N offspring individuals using crossover and mutation. Similarly, generate a choice key for each offspring individual to sample a one-path subnet from the supernet. And then use client sampling technique to download sampled subnets (download the choice keys from the second generation) for training and uploading the trained subnets to the server for aggregation.
- Download the supernet together with the choice keys of all parent and offspring individuals to all participating clients to evaluate the objectives. Upload all the objective values to the server and calculate the weighted average of the validation errors for each individual.
- Combine the parent and offspring individuals into a whole population. Perform environmental selection to select N new parents.
- Repeat Step 3 to Step 4 until the generation number reaches the pre-define maximum value.

Since only one path of the supernet needs to be trained on each client, this sampling approach can significantly reduce both upload and local computation costs. There is a small detail in Step 4 that downloads all the supernet model to every client to calculate the validation accuracies; thus, only choice keys are downloaded in the next generation since the whole supernet has been already downloaded in the last generation.

Online methods enable federated NAS systems to perform architecture search and train the model simultaneously. Both fitness evaluation thresholds, e.g. the number of communication rounds in federated learning and extra communication resources for training the searched models are not required using online approaches, which are highly desired for federated learning. However, the search space of online federated NAS is fairly limited, which affects the diversity of the architecture search.

Single- and multi-objective search

The aforementioned gradient-based federated NAS framework only considers and optimizes the model performances, which is usually not enough for federated learning, because federated NAS is naturally a multi-objective optimization problem. In addition to the maximization of the model performance, the payload (communication costs) to be transferred between the server and clients should be minimized. Single-objective optimization often aggregates conflict objectives into one objective using hyperparameters, while Pareto approaches aim to obtain a set of models presenting trade-off relationships between the conflicting objectives.

For example in RT-FedEvoNAS, the validation accuracy, model size and model floating point operations per second (FLOPs) of the sampled subnets are considered as the objec-

tives need to be optimized and NSGA-II [97] is used as the basic search algorithm. Finally, after several generations of evolutionary optimization, multiple well trained subnets can be obtained chosen from the trade-off solutions based on the user's preferences.

Open challenges

Currently, research on federated NAS is still very preliminary and several challenges remain to be solved.

Horizontally partitioned data

There is no general solution that can well solve the non-IID learning degradation problems in horizontal federated learning, let alone in federated NAS. The earliest work to explore non-IID data effect is proposed in [13]. The authors analyze the possible reasons for divergence in global model training on non-IID data and propose a strategy to mitigate this influence by globally sharing a small part of the data across all connected clients. However, this kind of data sharing intrinsically violate the scope of privacy-preserving scheme.

The federated distillation approaches [123,124] also have the potential risk of local data leakage. For the distillation, the teacher models are evaluated on mini-batches of unlabeled data on the server and their logits for mini-batch are used to train the student model on the server. The server can get a lot of local data information even on fake mini-batch data generated by local GAN generators [125].

Some statistical aggregation methods [14,126] are proposed to replace the original federated averaging algorithm (FedAvg). Both mathematical and experimental results prove that the proposed aggregation algorithms outperform the FedAvg on non-IID data. However, these approaches are often limited to some specific models and datasets, and it is unclear if they can show better performance for federated NAS frameworks. Hsieh et al. discuss the effect of non-IID data for DNNs in detail and different federated optimization methods are used upon different DNNs, such as GoogleNet [127], and ResNet [74]. Experiment results show that batch normalization [128] performs really poorly on non-IID data, but batch renormalization [129] and group normalization [130] are much more robust for non-IID data, which are much more suited for federated learning. Most recently, it is shown in [10] that ternary quantization is helpful in alleviating model divergence in federated learning, although its effectiveness remains to be validated on federated NAS.

Vertically partitioned data

Current federated NAS methods are all based on horizontal federated learning. Unlike horizontal federated learning, it

is really hard to determine whether the data are IID or non-IID in vertical federated learning since they are 'partitioned' towards the feature space.

Most existing vertical federated learning frameworks are built on two-party systems using simple linear models. Since only one party can hold the labels, the loss needs to be calculated on ciphertext; otherwise, the label information will be revealed. Then the gradients are very hard to calculate since the total loss is encrypted. Some approximation techniques like Taylor expansions [23,25] are often used to simplify the gradient calculations, which, however, may introduce strong biases for complex models like DNNs.

Overall, vertical federated NAS is totally different from horizontal federated NAS, which is in general still an unexplored research area.

Adversarial federated neural architecture search

Adversarial federated learning has two purposes: (1) inference of the client data information; (2) attack the global model to conduct backdoor [131] elements or even let the model unusable. And the adversary in federated learning can be one of participating clients or the central server, since we often assume the server is honest-but-curious. Thus, the server should also be considered as a potential risk.

Federated learning is still fragile to white box attacks since the model gradients and parameters still contain local data information. Geiping et al. [132] showed that local images can be reconstructed from the knowledge of model parameters (or gradients) by inverting gradients techniques. In addition, an adversarial GAN [125] generator can be developed on either the server [133] or the client side [134]. The adversary can reconstruct other participating clients' private data even if it has no knowledge of the label information.

Enthoven and Al-Ars [135] summarize most defence strategies used in federated learning, which can be categorized into three types: (1) subsample or compress the communicated gradients [6,7,136]; (2) differential privacy and SMC [19], and (3) robustness aggregation [137] using, e.g., the Byzantine resilient aggregation rule [138,139].

In general, finding robust model architectures in federated learning to defend against the adversarial attacks is still a hard task.

Encrypted federated neural architecture search

Homomorphic encryption technologies are often applied to prevent privacy leakage from the gradient information sent to the server. However, using homomorphic encryption in federated NAS system has two main difficulties.

First, homomorphic encryption, including encoding and encryption, is computationally expensive in federated learning. At first, all communicated model parameters need to

be encoded into large integers, because homomorphic encryption does not work on real numbers. Then the encoded parameters need to do modulus calculations with large prime numbers one by one. Unfortunately, modern deep neural network models contain millions of parameters, making the encryption process computationally extremely intensive. Therefore, developing a light weighted encryption method is an important yet challenging task for federated learning, let alone for federated NAS.

Second, the original federated encryption is introduced in which the server holds the public key and clients hold the private key. This framework is unsafe, because only one of clients uploads its secret key to the server. Therefore, a more advanced SMC approach is adopted that divides the whole secret key into several shards, and the server cannot decrypt the gradients unless it collected t (secret key recover threshold) key shards. Unfortunately, the encrypted gradients must be frequently transferred between the server and clients, which significantly increases the communication costs, since the local clients can only partially decrypt the gradients through their key shards. This is a big burden to the communication resources, which needs to be solved in the future.

Conclusion

In this survey paper, a brief overview of federated learning and NAS is provided, and a combination of both techniques, i.e., federated NAS is introduced. Given several remaining challenges in both federated learning and NAS, federated NAS becomes extremely challenging since many techniques developed in centralized NAS are no longer suited for federated NAS, and NAS will be subject to more constraints introduced by the federated learning environment. Two approaches to federated NAS are discussed, one is offline optimization and the other is online optimization. It is noted that offline evolutionary NAS methods are not applicable for many real-world scenarios, mainly because the offline approach performs architecture search and weight training separately and requires a large amount of communication costs. In addition, the performance of neural network under optimization must be acceptable for application and serious performance drop is not allowed. RT-FedEvoNAS [120] offers a solution to the above challenges, although its search space is highly constrained.

Despite that many grand challenges remain to be solved, federated NAS is of paramount practical significance for many real-world problems, where handcrafted deep neural networks may fail to work properly. We hope that this survey will help understand the promises and challenges in federated NAS, thereby triggering more research interests in developing new theories and algorithms, thereby promoting the

application of AI techniques to a wider range of fields where data privacy and security is a primary concern.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Elsken T, Metzen JH, Hutter F (2018) Neural architecture search: a survey, arXiv preprint [arXiv:1808.05377](https://arxiv.org/abs/1808.05377)
2. Yang Q, Liu Y, Chen T, Tong Y (2019) Federated machine learning: concept and applications. *ACM Trans Intell Syst Technol (TIST)* 10(2):1–19
3. McMahan B, Moore E, Ramage D, Hampson S, Arcas BAy (2017) Communication-efficient learning of deep networks from decentralized data. *Artif Intell Stat* 1273–1282
4. Zhang G-D, Zhao S-Y, Gao H, Li W-J (2018) Feature-distributed svrg for high-dimensional linear classification, arXiv preprint [arXiv:1802.03604](https://arxiv.org/abs/1802.03604)
5. McDonald R, Hall K, Mann G (2010) Distributed training strategies for the structured perceptron. In: *Human language technologies: The 2010 annual conference of the North American chapter of the association for computational linguistics*, 456–464
6. Shokri R, Shmatikov V (2015) Privacy-preserving deep learning. In: *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, 1310–1321
7. Konečný J, McMahan HB, Yu FX, Richtárik P, Suresh AT, Bacon D (2016) Federated learning: Strategies for improving communication efficiency. arXiv preprint [arXiv:1610.05492](https://arxiv.org/abs/1610.05492)
8. Caldas S, Konečný J, McMahan HB, Talwalkar A (2018) Expanding the reach of federated learning by reducing client resource requirements. arXiv preprint [arXiv:1812.07210](https://arxiv.org/abs/1812.07210)
9. Han S, Mao H, Dally WJ (2016) Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding. In: *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2–4, 2016, Conference Track Proceedings*, Y Bengio and Y LeCun, Eds., [Online]. [arXiv:1510.00149](https://arxiv.org/abs/1510.00149)
10. Xu J, Du W, Jin Y, He W, Cheng R (2020) Ternary compression for communication-efficient federated learning. *IEEE Trans Neural Netw Learn Syst*
11. Chen Y, Sun X, Jin Y (2019) Communication-efficient federated deep learning with layerwise asynchronous model update and temporally weighted aggregation. *IEEE Trans Neural Netw Learn Syst*

12. Zhu H, Jin Y (2020) Multi-objective evolutionary federated learning. *IEEE Trans Neural Netw Learn Syst* 31(4):1310–1322
13. Zhao Y, Li M, Lai L, Suda N, Civin D, Chandra V (2018) Federated learning with non-iid data. arXiv preprint [arXiv:1806.00582](https://arxiv.org/abs/1806.00582)
14. Li T, Sahu AK, Zaheer M, Sanjabi M, Talwalkar A, Smith V (2018) Federated optimization in heterogeneous networks. arXiv preprint [arXiv:1812.06127](https://arxiv.org/abs/1812.06127)
15. Ji S, Pan S, Long G, Li X, Jiang J, Huang Z (2019) Learning private neural language modeling with attentive aggregation. In: 2019 International joint conference on neural networks (IJCNN). IEEE, 1–8
16. Phong LT, Aono Y, Hayashi T, Wang L, Moriai S (2018) Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Trans Inf Forensics Secur* 13(5):1333–1345
17. Armknecht F, Boyd C, Carr C, Gjosteen K, Jäschke A, Reuter CA, Strand M (2015) A guide to fully homomorphic encryption. Cryptology ePrint Archive, Report 2015/1192, <https://eprint.iacr.org/2015/1192>
18. Lindell Y (2005) Secure multiparty computation for privacy preserving data mining. In: *Encyclopedia of Data Warehousing and Mining*. IGI Global, 1005–1009
19. Bonawitz K, Ivanov V, Kreuter B, Marcedone A, McMahan HB, Patel S, Ramage D, Segal A, Seth K (2017) Practical secure aggregation for privacy preserving machine learning. Cryptology ePrint Archive, Report 2017/281, <https://eprint.iacr.org/2017/281>
20. Dwork C (2008) Differential privacy: a survey of results. In: *International conference on theory and applications of models of computation*. Springer, 1–19
21. Geyer RC, Klein T, Nabi M (2017) Differentially private federated learning: a client level perspective. arXiv preprint [arXiv:1712.07557](https://arxiv.org/abs/1712.07557)
22. Abadi M, Chu A, Goodfellow I, McMahan HB, Mironov I, Talwar K, Zhang L (2016) Deep learning with differential privacy. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318
23. Hardy S, Henecka W, Ivey-Law H, Nock R, Patrini G, Smith G, Thorne B (2017) Private federated learning on vertically partitioned data via entity resolution and additively homomorphic encryption. arXiv preprint [arXiv:1711.10677](https://arxiv.org/abs/1711.10677)
24. Yang S, Ren B, Zhou X, Liu L (2019) Parallel distributed logistic regression for vertical federated learning without third-party coordinator. arXiv preprint [arXiv:1911.09824](https://arxiv.org/abs/1911.09824)
25. Liu Y, Kang Y, Xing C, Chen T, Yang Q (2020) A secure federated transfer learning framework. *IEEE Intell Syst*
26. Feng S, Yu H (2020) Multi-participant multi-class vertical federated learning. arXiv preprint [arXiv:2001.11154](https://arxiv.org/abs/2001.11154)
27. Liu Y, Zhang X, Wang L (2020) Asymmetrically vertical federated learning. arXiv preprint [arXiv:2004.07427](https://arxiv.org/abs/2004.07427)
28. Freedman MJ, Nissim K, Pinkas B (2004) Efficient private matching and set intersection. In: Cachin C, Camenisch JL (eds) *Advances in cryptology - EUROCRYPT 2004*. Springer, Berlin Heidelberg, pp 1–19
29. Liang G, Chawathe SS (2004) Privacy-preserving inter-database operations. In: Chen H, Moore R, Zeng DD, Leavitt J (eds) *Intelligence and Security Informatics*. Springer, Berlin Heidelberg, pp 66–82
30. Bergstra J, Bengio Y (2012) Random search for hyper-parameter optimization. *J Mach Learn Res* 13(1):281–305
31. Li L, Talwalkar A (2020) Random search and reproducibility for neural architecture search. In: *Uncertainty in Artificial Intelligence*. PMLR, 367–377
32. Rasmussen CE (2003) *Gaussian processes in machine learning*. Summer School on Machine Learning. Springer, New York, pp 63–71
33. Swersky K, Snoek J, Adams RP (2014) Freeze-thaw bayesian optimization. arXiv preprint [arXiv:1406.3896](https://arxiv.org/abs/1406.3896)
34. Zheng X, Ji R, Tang L, Zhang B, Liu J, Tian Q (2019) Multinomial distribution learning for effective neural architecture search. In: *Proceedings of the IEEE International Conference on Computer Vision*, 1304–1313
35. Bottou L (2012) *Stochastic gradient descent tricks*. Neural networks: tricks of the trade. Springer, New York, pp 421–436
36. Zoph B, Le QV (2017) Neural architecture search with reinforcement learning
37. Zoph B, Yuret D, May J, Knight K (2016) Transfer learning for low-resource neural machine translation. arXiv preprint [arXiv:1604.02201](https://arxiv.org/abs/1604.02201)
38. Zoph B, Vasudevan V, Shlens J, Le QV (2018) Learning transferable architectures for scalable image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8697–8710
39. Xie L, Yuille A (2017) Genetic cnn. In: *Proceedings of the IEEE international conference on computer vision*, 1379–1388
40. Sun Y, Xue B, Zhang M, Yen GG (2019) Completely automated cnn architecture design based on blocks. *IEEE Trans Neural Netw Learn Syst* 31(4):1242–1254
41. Sun Y, Xue B, Zhang M, Yen GG, Lv J (2020) Automatically designing cnn architectures using the genetic algorithm for image classification. *IEEE Trans Cybern*
42. Zela A, Klein A, Falkner S, Hutter F (2018) Towards automated deep learning: Efficient joint neural architecture and hyperparameter search. arXiv preprint [arXiv:1807.06906](https://arxiv.org/abs/1807.06906)
43. Sun Y, Wang H, Xue B, Jin Y, Yen GG, Zhang M (2019) Surrogate-assisted evolutionary deep learning using an end-to-end random forest-based performance predictor. *IEEE Trans Evol Comput* 24(2):350–364
44. Zhang H, Kiranyaz S, Gabbouj M (2018) Finding better topologies for deep convolutional neural networks by evolution. arXiv preprint [arXiv:1809.03242](https://arxiv.org/abs/1809.03242)
45. Jin H, Song Q, Hu X (2018) Efficient neural architecture search with network morphism, vol. 9. arXiv preprint [arXiv:1806.10282](https://arxiv.org/abs/1806.10282)
46. Pham H, Guan MY, Zoph B, Le QV, Dean J (2018) Efficient neural architecture search via parameter sharing. arXiv preprint [arXiv:1802.03268](https://arxiv.org/abs/1802.03268)
47. Guo Z, Zhang X, Mu H, Heng W, Liu Z, Wei Y, Sun J (2020) Single path one-shot neural architecture search with uniform sampling. In: *Proceedings of the European Conference on Computer Vision (ECCV)*
48. Dong X, Yang Y (2019) One-shot neural architecture search via self-evaluated template network. In: *Proceedings of the IEEE International Conference on Computer Vision*, 3681–3690
49. Li X, Lin C, Li C, Sun M, Wu W, Yan J, Ouyang W (2020) Improving one-shot NAS by suppressing the posterior fading. *Proc IEEE/CVF Conf Comput Vis Pattern Recognit* 13(836–13):845
50. You S, Huang T, Yang M, Wang F, Qian C, Zhang C (2020) GreedyNAS: Towards fast one-shot NAS with greedy supernet. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1999–2008*
51. Watkins CJH (1989) Learning from delayed rewards
52. Baker B, Gupta O, Naik N, Raskar R (2016) Designing neural network architectures using reinforcement learning. arXiv preprint [arXiv:1611.02167](https://arxiv.org/abs/1611.02167)
53. Zhong Z, Yang Z, Deng B, Yan J, Wu W, Shao J, Liu C-L (2020) Blockqnn: Efficient block-wise neural network architecture generation. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*
54. Pasunuru R, Bansal M (2019) Continual and multi-task architecture search. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, 1911–1922*

55. Bender G, Liu H, Chen B, Chu G, Cheng S, Kindermans P-J, Le QV (2020) Can weight sharing outperform random architecture search? an investigation with TuNAS. *Proc IEEE/CVF Conf Comput Vis Pattern Recognit* 14(323–14):332
56. Liu Y, Jia X, Tan M, Vemulapalli R, Zhu Y, Green B, Wang X (2020) Search to distill: Pearls are everywhere but not the eyes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7539–7548
57. Hinton G, Vinyals O, Dean J (2015) Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*
58. Back T (1996) *Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms*. Oxford University Press, Oxford
59. Banzhaf W, Nordin P, Keller RE, Francone FD (1998) *Genetic programming—an introduction: On the automatic evolution of computer programs and its applications*, dpunkt. verlag and morgan kaufmann publishers Inc., San Francisco, California
60. Schmitt LM (2001) Theory of genetic algorithms. *Theoret Comput Sci* 259(1–2):1–61
61. Yao X (1999) Evolving artificial neural networks. *Proc IEEE* 87(9):1423–1447
62. Floreano D, Dürr P, Mattiussi C (2008) Neuroevolution: from architectures to learning. *Evol Intel* 1(1):47–62
63. Stanley KO, D'Ambrosio DB, Gauci J (2009) A hypercube-based encoding for evolving large-scale neural networks. *Artif life* 15(2):185–212
64. Jozefowicz R, Zaremba W, Sutskever I (2015) An empirical exploration of recurrent network architectures. In: *International conference on machine learning*, 2342–2350
65. Stanley KO, Miikkulainen R (2002) Evolving neural networks through augmenting topologies. *Evol Comput* 10(2):99–127
66. Galván E, Mooney P (2020) Neuroevolution in deep neural networks: Current trends and future challenges. *arXiv preprint arXiv:2006.05415*
67. Darwish A, Hassanien AE, Das S (2020) A survey of swarm and evolutionary computing approaches for deep learning. *Artif Intell Rev* 53(3):1767–1812
68. Wang H, Jiao L, Yao X (2014) Two_arch2: an improved two-archival algorithm for many-objective optimization. *IEEE Trans Evol Comput* 19(4):524–541
69. Sun Y, Yen GG, Yi Z (2018) Igd indicator-based evolutionary algorithm for many-objective optimization problems. *IEEE Trans Evol Comput* 23(2):173–187
70. Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M et al (2015) Imagenet large scale visual recognition challenge. *Int J Comput Vis* 115(3):211–252
71. Miikkulainen R, Liang J, Meyerson E, Rawal A, Fink D, Francon O, Raju B, Shahrzad H, Navruzyan A, Duffy N et al (2019) Evolving deep neural networks. In: *Artificial Intelligence in the Age of Neural Networks and Brain Computing*. Elsevier 293–312
72. Liang J, Meyerson E, Miikkulainen R (2018) Evolutionary architecture search for deep multitask networks. In: *Proceedings of the Genetic and Evolutionary Computation Conference*, 466–473
73. Real E, Moore S, Selle A, Saxena S, Suematsu YL, Tan J, Le QV, Kurakin A (2017) Large-scale evolution of image classifiers. In: *International Conference on Machine Learning*, 2902–2911
74. He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778
75. Huang G, Liu Z, Van Der Maaten L, Weinberger KQ (2017) Densely connected convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4700–4708
76. Irwin-Harris W, Sun Y, Xue B, Zhang M (2019) A graph-based encoding for evolutionary convolutional neural network architecture design. In: *2019 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 546–553
77. Suganuma M, Shirakawa S, Nagao T (2017) A genetic programming approach to designing convolutional neural network architectures. In: *Proceedings of the genetic and evolutionary computation conference*, 497–504
78. Harding S (2008) Evolution of image filters on graphics processor units using cartesian genetic programming. In: *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*. IEEE, 1921–1928
79. Miller JF, Smith SL (2006) Redundancy and computational efficiency in cartesian genetic programming. *IEEE Trans Evol Comput* 10(2):167–174
80. Sun Y, Xue B, Zhang M, Yen GG (2019) Evolving deep convolutional neural networks for image classification. *IEEE Trans Evol Comput* 24(2):394–407
81. Rasmussen CE, Ghahramani Z (2001) Occam's razor. In: *Advances in neural information processing systems*, 294–300
82. Hawkins DM (2004) The problem of overfitting. *J Chem Inf Comput Sci* 44(1):1–12
83. Liu H, Simonyan K, Vinyals O, Fernando C, Kavukcuoglu K (2018) Hierarchical representations for efficient architecture search. In: *International Conference on Learning Representations*
84. Lu Z, Whalen I, Boddeti V, Dhebar Y, Deb K, Goodman E, Banzhaf W (2019) Nsga-net: neural architecture search using multi-objective genetic algorithm. In: *Proceedings of the Genetic and Evolutionary Computation Conference*, 419–427
85. E. Real, A. Aggarwal, Y. Huang, and Q. V. Le, Regularized evolution for image classifier architecture search. In: *Proceedings of the aaai conference on artificial intelligence*, vol. 33, 2019, pp. 4780–4789
86. Zhou D, Zhou X, Zhang W, Loy CC, Yi S, Zhang X, Ouyang W (2020) EcoNAS: Finding proxies for economical neural architecture search. *Proc IEEE/CVF Conf Comput Vis Pattern Recognit* 11(396–11):404
87. Lu Z, Whalen I, Dhebar Y, Deb K, Goodman E, Banzhaf W, Boddeti VN (2019) Multi-criterion evolutionary design of deep convolutional neural networks. *arXiv preprint arXiv:1912.01369*
88. Jin Y (2011) Surrogate-assisted evolutionary computation: recent advances and future challenges. *Swarm Evol Comput* 1(2):61–70
89. Broomhead DS, Lowe D (1988) Radial basis functions, multi-variable functional interpolation and adaptive networks. Royal Signals and Radar Establishment Malvern (United Kingdom). Tech. Rep
90. Baker B, Gupta O, Raskar R, Naik N (2017) Accelerating neural architecture search using performance prediction. *arXiv preprint arXiv:1705.10823*
91. Jeong S, Murayama M, Yamamoto K (2005) Efficient optimization design method using kriging model. *J Aircr* 42(2):413–420
92. Dai X, Zhang P, Wu B, Yin H, Sun F, Wang Y, Dukhan M, Hu Y, Wu Y, Jia Y et al (2019) Chamnet: Towards efficient network design through platform-aware model adaptation, 11 398–11 407
93. Wang H, Jin Y, Doherty J (2017) Committee-based active learning for surrogate-assisted particle swarm optimization of expensive problems. *IEEE Trans Cybern* 47(9):2664–2677
94. Shahriari B, Swersky K, Wang Z, Adams RP, De Freitas N (2015) Taking the human out of the loop: a review of bayesian optimization. *Proc IEEE* 104(1):148–175
95. Lu Z, Deb K, Goodman E, Banzhaf W, Boddeti VN (2020) Nsganetv2: Evolutionary multi-objective surrogate-assisted neural architecture search. In: *Proceedings of the European Conference on Computer Vision (ECCV)*
96. Kim Y-H, Reddy B, Yun S, Seo C (2017) Nemo: Neuro-evolution with multiobjective optimization of deep neural network for speed and accuracy. In: *ICML 2017 AutoML Workshop*

97. Deb K, Pratap A, Agarwal S, Meyarivan T (2002) A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Trans Evol Comput* 6(2):182–197
98. Dong J-D, Cheng A-C, Juan D-C, Wei W, Sun M (2018) Dppnet: Device-aware progressive search for pareto-optimal neural architectures. In: *Proceedings of the European Conference on Computer Vision (ECCV)*, 517–531
99. Liu C, Zoph B, Neumann M, Shlens J, Hua W, Li L-J, Fei-Fei L, Yuille A, Huang J, Murphy K (2018) Progressive neural architecture search. In: *Proceedings of the European Conference on Computer Vision (ECCV)*, 19–34
100. Elsken T, Metzen JH, Hutter F (2018) Efficient multi-objective neural architecture search via lamarckian evolution. *arXiv preprint arXiv:1804.09081*
101. Wei T, Wang C, Rui Y, Chen CW (2016) Network morphism. In: *International Conference on Machine Learning*, 564–572
102. Jin Y, Sendhoff B (2008) Pareto-based multiobjective machine learning: an overview and case studies. *IEEE Trans Syst Man Cybern Part C (Applications and Reviews)* 38(3):397–415
103. Lu Z, Sreekumar G, Goodman E, Banzhaf W, Deb K, Bodeti VN (2020) Neural architecture transfer. *arXiv preprint arXiv:2005.05859*
104. Shin R, Packer C, Song D (2018) Differentiable neural network architecture search
105. Ahmed K, Torresani L (2018) askconnect: Connectivity learning by gradient descent. In: *Proceedings of the European Conference on Computer Vision (ECCV)*, 349–365
106. Fang J, Sun Y, Zhang Q, Li Y, Liu W, Wang X (2020) Densely connected search space for more flexible neural architecture search. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10 628–10 637
107. Cai H, Zhu L, Han S (2018) ProxylessNAS: Direct neural architecture search on target task and hardware. In: *International Conference on Learning Representations*
108. Lorraine J, Vicol P, Duvenaud D (2020) Optimizing millions of hyperparameters by implicit differentiation. In: *International Conference on Artificial Intelligence and Statistics*, 1540–1552
109. Liu H, Simonyan K, Yang Y (2018) Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*
110. Dong X, Yang Y (2019) Searching for a robust neural architecture in four gpu hours. In: *Proceedings of the IEEE Conference on computer vision and pattern recognition*, 1761–1770
111. Xie S, Zheng H, Liu C, Lin L (2018) SNAS: stochastic neural architecture search. In: *International Conference on Learning Representations*
112. Li G, Qian G, Delgadillo IC, Muller M, Thabet A, Ghanem B (2020) Sgas: Sequential greedy architecture search. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1620–1630
113. Xu Y, Xie L, Zhang X, Chen X, Qi G-J, Tian Q, Xiong H (2019) Pc-darts: Partial channel connections for memory-efficient architecture search. In: *International Conference on Learning Representations*
114. Gao C, Chen Y, Liu S, Tan Z, Yan S (2020) AdversarialNas: Adversarial neural architecture search for GANs. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5680–5689
115. Wan A, Dai X, Zhang P, He Z, Tian Y, Xie S, Wu B, Yu M, Xu T, Chen K et al (2020) Fbnetv2: differentiable neural architecture search for spatial and channel dimensions. *Proc IEEE/CVF Conf Comput Vis Pattern Recognit* 12(965–12):974
116. He C, Ye H, Shen L, Zhang T (2020) Milenas: efficient neural architecture search via mixed-level reformulation. *Proc IEEE/CVF Conf Comput Vis Pattern Recognit* 11(993–12):002
117. Xu M, Zhao Y, Bian K, Huang G, Mei Q, Liu X (2020) Neural architecture search over decentralized data. *arXiv preprint arXiv:2002.06352*
118. Luo J-H, Wu J, Lin W (2017) Thinet: A filter level pruning method for deep neural network compression. In: *Proceedings of the IEEE international conference on computer vision*, 5058–5066
119. He C, Annavaram M, Avestimehr S (2020) Fednas: Federated deep learning via neural architecture search. *arXiv preprint arXiv:2004.08546*
120. Zhu H, Jin Y (2020) Real-time federated evolutionary neural architecture search. *arXiv preprint arXiv:2003.02793*
121. Trick MA (1992) A linear relaxation heuristic for the generalized assignment problem. *Naval Res Logist (NRL)* 39(2):137–151
122. Bender G, Kindermans P-J, Zoph B, Vasudevan V, Le Q (2018) Understanding and simplifying one-shot architecture search. In: *International Conference on Machine Learning*, 550–559
123. Jeong E, Oh S, Kim H, Park J, Bennis M, Kim S-L (2018) Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data. *arXiv preprint arXiv:1811.11479*
124. Lin T, Kong L, Stich SU, Jaggi M (2020) Ensemble distillation for robust model fusion in federated learning. *arXiv preprint arXiv:2006.07242*
125. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial nets. In: *Advances in neural information processing systems*, 2672–2680
126. Mohri M, Sivek G, Suresh AT (2019) Agnostic federated learning. *arXiv preprint arXiv:1902.00146*
127. Szegedy C, Liu W, Jia Y, Sermanet P, Reed S, Anguelov D, Erhan D, Vanhoucke V, Rabinovich A (2015) Going deeper with convolutions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1–9
128. Ioffe S, Szegedy C (2015) Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*
129. Ioffe S (2017) Batch renormalization: Towards reducing mini-batch dependence in batch-normalized models. In: *Advances in neural information processing systems*, 1945–1953
130. Wu Y, He K (2018) Group normalization. In: *Proceedings of the European conference on computer vision (ECCV)*, 3–19
131. Bagdasaryan E, Veit A, Hua Y, Estrin D, Shmatikov V (2020) How to backdoor federated learning. ser. *Proceedings of Machine Learning Research*, S. Chiappa and R. Calandra, Eds., vol. 108. Online: PMLR, 26–28 Aug 2020, pp. 2938–2948. [Online]. <http://proceedings.mlr.press/v108/bagdasaryan20a.html>
132. Geiping J, Bauermeister H, Dröge H, Moeller M (2020) Inverting gradients—how easy is it to break privacy in federated learning? *arXiv preprint arXiv:2003.14053*
133. Wang Z, Song M, Zhang Z, Song Y, Wang Q, Qi H (2019) Beyond inferring class representatives: User-level privacy leakage from federated learning. In: *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2512–2520
134. Hitaj B, Ateniese G, Perez-Cruz F (2017) Deep models under the gan: information leakage from collaborative deep learning. In: *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 603–618
135. Enthoven D, Al-Ars Z (2020) An overview of federated deep learning privacy attacks and defensive strategies. *arXiv preprint arXiv:2004.04676*

136. Yoon J, Jeong W, Lee G, Yang E, Hwang SJ (2020) Federated continual learning with adaptive parameter communication. arXiv preprint [arXiv:2003.03196](https://arxiv.org/abs/2003.03196)
137. Pillutla K, Kakade SM, Harchaoui Z (2019) Robust aggregation for federated learning. arXiv preprint [arXiv:1912.13445](https://arxiv.org/abs/1912.13445)
138. Lamport L, Shostak R, Pease M (2019) The byzantine generals problem. In: Concurrency: the Works of Leslie Lamport, 203–226
139. Mhamdi EME, Guerraoui R, Rouault S (2018) The hidden vulnerability of distributed learning in byzantium. arXiv preprint [arXiv:1802.07927](https://arxiv.org/abs/1802.07927)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.