

Communication-Efficient and Distributed Learning Over Wireless Networks: Principles and Applications

†Jihong Park, Sumudu Samarakoon, Anis Elgabri, ‡Joongheon Kim, Mehdi Bennis, *Seong-Lyun Kim, and **Mérrouane Debbah

Abstract—Machine learning (ML) is a promising enabler for the fifth generation (5G) communication systems and beyond. By imbuing intelligence into the network edge, edge nodes can proactively carry out decision-making, and thereby react to local environmental changes and disturbances while experiencing zero communication latency. To achieve this goal, it is essential to cater for high ML inference accuracy at scale under time-varying channel and network dynamics, by continuously exchanging fresh data and ML model updates in a distributed way. Taming this new kind of data traffic boils down to improving the communication efficiency of distributed learning by optimizing communication payload types, transmission techniques, and scheduling, as well as ML architectures, algorithms, and data processing methods. To this end, this article aims to provide a holistic overview of relevant communication and ML principles, and thereby present communication-efficient and distributed learning frameworks with selected use cases.

1 SIGNIFICANCE AND MOTIVATION

The pursuit of extremely stringent latency and reliability guarantees is essential in the fifth generation (5G) communication system and beyond [1], [2]. In a wirelessly automated factory, the remote control of assembly robots should provision the same level of target latency and reliability offered by existing wired factory systems. To this end, for instance, control packets should be delivered within 1 ms with 99.99999% reliability [3]–[5]. Things are becoming even more challenging in the emerging mission-critical applications beyond 5G. A prime example is the forthcoming non-terrestrial networks consisting of a massive constellation of low-altitude earth orbit (LEO) satellites [6]–[11]. Given such

a large number of satellites orbiting with the speed over 8 km per second, emergency control is no longer an option for avoiding collisions with other satellites and space debris.

Unfortunately, traditional remote control methods are unfit for such safety operations. These methods often postulate known stationary channel and network topological models while focusing primarily on maximizing data rates. Such model-based and best-effort solutions are far from enough to meet the challenging latency and reliability requirements under limited radio resources and randomness on wireless channels and network topologies in practice.

Realizing the aforementioned pressing concern has recently sparked huge attention to the introduction of machine learning (ML) based approaches into communication system designs [12]–[16]. By leveraging ML at the network edge, each edge node can proactively carry out decision-making based on its local predictions, thereby experiencing zero latency [2], [17]. Furthermore, real data observations construct these ML models that directly reflect the environment in reality without modeling artifacts. In these respects, one may misapprehend that communication becomes less important in 5G and beyond where everything is locally predictable. The answer is the opposite, as accurate ML prediction cannot be achieved and sustained without communication.

In fact, every edge node can only observe a tiny fraction of the entire environment, in terms of time, space, and extreme events. Collectively utilizing these dispersed local data is thus a prerequisite to train and run data-driven ML models with high prediction or inference accuracy. For this purpose, local data do not have to be directly exchanged across edge nodes, which may violate data privacy, not to mention the non-negligible latency. Alternatively, by leveraging federated learning (FL), it is possible to exchange ML model parameters that reflect the accumulated data observed by each ML model without revealing raw data [18]. Similarly, one can exchange ML model outputs [19], [20] or hidden activations [21] for higher communication efficiency while preserving data privacy. Such a communication is not a one-time event, since a trained ML model can easily be outdated and should thus be continually re-trained under time-varying data distributions and environments. As a consequence, ML will not only be a key enabler of future communication systems, but also be one major source of data traffic, which warrants taming the new kind of traffic

†J. Park is with the School of Information Technology, Deakin University, VIC 3220, Geelong, Australia (email: jihong.park@deakin.edu.au).

S. Samarakoon, A. Elgabri, and M. Bennis are with the Centre for Wireless Communications, University of Oulu, Oulu 90014, Finland (email: {sumudu.samarakoon, anis.elgabri, mehdi.bennis}@oulu.fi).

J. Kim is with the School of Electrical Engineering, Korea University, Seoul 02841, Korea (e-mail: joongheon@korea.ac.kr).

*S.-L. Kim is with the School of Electrical and Electronic Engineering, Yonsei University, Seoul 03722, Korea (email: slkim@yonsei.ac.kr)

**M. Debbah is with Université Paris-Saclay, CNRS, CentraleSupélec, Gif-sur-Yvette 91190, France (email: merouane.debbah@centralesupelec.fr) and the Lagrange Mathematical and Computing Research Center, Paris 75007, France. This work was supported in part by the INFOTECH Project NOOR, by the NEGEIN project, by the EU-CHISTERA projects LeadingEdge and CON-NECT, by the EU-H2020 project (IntelloT) and the Academy of Finland projects MISSION and SMARTER, and by the Institute for Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2018-0-00170, Virtual Presence in Moving Objects through 5G).




Fig. 1. An overview of this article that aims to present communication-efficient and distributed learning frameworks (Sec. 6) by applying advanced machine learning (ML) and communication principles (Sec. 4 and 5) to baseline distributed learning algorithms (Sec. 3) for addressing key ML and communication challenges (Sec. 2).

generated by distributed learning. Furthermore, communication environments have a considerable impact on the performance of ML. Indeed, temporal network topology variations and uplink-downlink channel asymmetry determine learning stragglers. In analog transmissions, channel fluctuations directly distort communicating information [22], affecting the ML accuracy and data privacy. This mandates to co-design distributed ML and communication operations.

Spurred by the aforementioned motivations, this article aims to present communication-efficient and distributed learning frameworks built upon jointly optimizing the types of communication payloads, transmissions, and scheduling as well as ML architectures and algorithms under wireless channel dynamics and network topology variations. To reach the overarching goal, as visualized in Fig. 1, this article is structured as follows. In Sec. 2, major technical challenges are summarized. In Sec. 3, existing communication-efficient and distributed learning frameworks and their limitations are briefly reviewed. To improve these vanilla distributed learning frameworks, several ML and communication design principles are discussed in Sec. 4 and 5. Finally, selected applications of such principles and their effectiveness are elaborated in Sec. 6, followed by concluding remarks in Sec. 7.

Note that there has been a recent upsurge of FL and relevant distributed learning frameworks [23], [24], which have been extensively studied in both ML and wireless communication communities. The fundamentals, key challenges, and recent advances of FL have been well summarized in [25]–[27]. With wireless connectivity, distributed learning frameworks have been overviewed in [13], [28], [29]. For

privacy and security issues, the challenges and opportunities of FL have been reviewed in [30]–[33]. A complete list of related works on distributed learning frameworks has been provided in [34].

As opposed to these works focusing mostly on high-level discussions often without demonstrating key applications, this article is driven primarily by use cases that are tightly linked to the underlying ML and communication principles, as visualized in Fig. 1. Compared to our preceding work [13], due to a similar way of presentation, a part of background discussions has been inevitably repeated for the sake of completeness. Nevertheless, in this article we have introduced new use cases reflecting more realistic communication and advanced ML techniques. They are associated with new communication and ML principles, underpinning the contributions of this work.

2 KEY CHALLENGES

Towards understanding the underlying black-box operations of ML, centralized ML architectures have been the prime focus in the theoretical studies. With the paradigm shift from cloud-centric to on-device ML, above theoretical analysis cannot be readily applicable to investigate the current distributed ML architecture. In this view, we identify several key challenges that needs to be addressed in designing distributed learning over wireless networks as discussed next.

Data Shortage. One of the main downsides of shifting from cloud to device in data-driven ML is the limited access to sufficiently large datasets. Devices with low exposure and storage may not be able to accumulate rich datasets, in

which training on-devices may lose generalization and are susceptible to unseen data [35]. Data acquisition within the device could result in higher end-to-end training latencies and/or outdated ML models in the presence of dynamic data. To overcome the data shortage, robust and collaborative ML designs need to be investigated.

Non-IID Data. User-generated data could be highly personalized (e.g. different angles and frame rates of surveillance cameras,), imbalanced (e.g., some labels corresponding to extreme events have much fewer samples than other labels), and multimodal (e.g., temperature and humidity sensors for weather prediction), all of which are described as non-independent and identically distributed (IID) data. Under non-IID data, it is common that the accuracy and convergence speed of distributed learning are significantly degraded [36]–[40]. Furthermore, majority of the analytical frameworks are devised on the training with IID data and cannot be easily extended towards distributed learning over non-IID data [27]. Hence, deriving convergence characteristic and reliability and robustness guarantees with on-device learning is a daunting task.

Data Privacy. Data owned by the devices may contain privacy-sensitive information, and thus, exchanging ML model parameters instead data is widely used in distributed learning. Yet, the exposed model parameters could be reversely traced, in which, privacy is only partially preserved [41]. To further enhance privacy, adopting extra coding, introducing noise to shared parameters, and exchanging redundant information are some viable solutions. However, each of above solutions introduce additional challenges (e.g., increased processing delays with extra coding, loss of inference accuracy due to the excess noise, and extra communication delays with redundant information).

Computing Resource Limitation. Training and operating ML models require huge computation processor energy, memory, and high-speed inter-processor communication links, which is commonly not available at battery-limited small edge devices. Hence, deep learning computations are often carried out at a cloud server using high-performance computing (HPC) resources [42] consisting of graphics processing units (GPUs), each of which is equipped with thousands of core processing units (e.g., NVIDIA GTX 2080 Ti has 4352 CUDA cores and 544 tensor processing units [43]). This cannot be expected to shift to the network edge without simplifying their complexities. In addition, due to the limited energy and memory/storage at edge devices, processing low-complex small models and tasks is of the utmost importance. In this view, designs of energy-efficient *low-precision* ML and *binary neural networks* (NNs) need to be considered with distributed learning [44]–[46].

Communication Resource Limitation. Relying on the limited bandwidth, multitude of devices and services share the same wireless resources. It is thus susceptible to high interference and intermittent connectivity, restraining the distributed learning accuracy and speed [13]. Mobile service operators are restricted to limited frequency bands as well as bandwidths, in which the difficulties on ensuring reliable and low-latency connectivity for training devices grow exponentially as the network scales. The problem can

be partly ameliorated by introducing more bandwidth to the network via the usage of high frequency bands such as millimeter waves (mmWaves). Nevertheless, the massive use of high frequency bands not only increases energy footprints and device temperature [47], but also suffers from their unreliable nature of channel conditions, in terms of high distance attenuation and sensitivity to blockages [48]. While increasing/optimizing transmit power and adopting encoding-decoding techniques can be beneficial in terms of enhancing reliable connectivity, training devices may not be able to exploit them with the limited power availability [49]. Therefore, communication resource management is a key aspect on realizing distributed learning.

Poor Channel Conditions. Distributed learning over large number of devices collaborating one another relies on the inter-device communication over wireless links. Under wireless channel dynamics, communication among devices is likely to be affected by poor channel conditions and the transmission noise yielding increased training latencies and losses in both training and inference accuracy [13]. With the limited wireless resources, it is crucial to adopt existing communication techniques (e.g., scheduling, coding, quantizing, relaying, interference managing, millimeter wave communication etc.) and to extend them considering the aspects of distributed learning (e.g., guarantees on training latency, accuracy, reliability, and robustness).

Time-Varying Network Topology. Mobility is inherited in devices, in which, distributed learning needs to cope with dynamic network topologies. With time-varying networks, learning agents are affected by loss of connectivity, inconsistent and asynchronous collaboration, frequent model mismatches, and tendency of having outdated data and models [35]. Developing distributed training mechanisms and analyzing them over above dynamics is extremely difficult. Resorting to predictive/proactive techniques and recasting the interactions among many agents to simplified statistical models are essential for learning over dynamic wireless network topologies.

3 RELATED DISTRIBUTED LEARNING METHODS

Distributed ML algorithms are briefly categorized into the methods exchanging model parameters, model outputs, and hidden activations, with or without the aid of a parameter server. In this section, we introduce representative distributed ML methods, followed by identifying the limitations of these vanilla approaches, calling for applying new key principles and developing advanced ML frameworks to be elaborated in the next sections.

3.1 Federated Learning (FL)

FL is a distributed training framework, which has been successfully adopted for Google’s predictive keyboards [50] and many other use cases in the areas of healthcare, intelligent transportation, and industrial automation [27], [51]. In essence, FL is designed to periodically upload workers’ model parameters (e.g., NN weights and/or gradients) during local training to a parameter server that performs model averaging and broadcasts the resultant global model

to all workers [52]. Here, avoiding raw data exchanges preserves data privacy, while adjusting the uploading period improves communication efficiency.

Recent studies have investigated different training aspects including personalization (i.e., multi-task learning) [53], robustness guarantees [54], [55], and training over dynamic topologies [56]. To further improve the data privacy against the attacks inverting model parameters into raw data [57], [58], various privacy-preserving methods have been investigated, such as injecting fine-tuned noise into model parameters via a differential privacy mechanism [31], [59]–[61] and mixing model parameters over the air via analog transmissions [62], [63]. Still, one critical issue of FL is that its communication overhead is proportional to the number of model parameters. Consequently, FL struggles with supporting deep NNs over capacity-limited wireless channels.

3.2 Group ADMM (GADMM)

The parameter server in FL cannot be connected with far-away workers. Furthermore, it is vulnerable to a single point of attack or failure [64]. In this regard, leveraging the alternating direction method of multipliers (ADMM) method, group ADMM (GADMM) aims to enable distributed learning without any central entity while communicating only with neighboring workers [65]. To this end, GADMM divides the workers into *head* and *tail* groups. Each worker from head or tail group exchanges variables with only two workers from the tail/head group forming a chain. At each iteration, every head worker first updates its primal variable (i.e., models) in parallel by minimizing the augmented Lagrangian function defined in ADMM, while utilizing its two neighboring tail workers' models in the previous iteration. Once head workers update their models, each worker transmits its updated model to its two neighbors from the tail group. Then, following the same way, every tail worker updates its model by utilizing its two neighboring head workers' models received in the current iteration. Finally, the dual variables are updated locally at each worker.

With GADMM, at every communication round, only half of the workers are competing for the limited bandwidth. Moreover, by limiting the communication only to the two neighboring workers, the communication energy can significantly be reduced. Nonetheless, GADMM relies on model parameter exchanges as in FL whose communication payload size increases with the number of parameters, limiting the scalability of GADMM particularly under deep NNs.

3.3 Federated Distillation (FD)

Modern deep NN architectures often have a large number of model parameters. For instance, GPT-3 model is a state-of-the-art NN architecture for natural language processing (NLP) tasks, and has 175 billion parameters corresponding to over 350 GB [66]. Exchanging the sheer amount of deep NN model parameter is costly, hindering frequent communications particularly under limited wireless resources. Alternatively, FD only exchanges the models' outputs whose dimensions are much smaller than the model sizes (e.g., 10 classes in the MNIST dataset). To illustrate, in a classification

task, each worker runs local iterations while storing the average model output (i.e., logit) per class. At a regular interval, these local average outputs are uploaded to the parameter server aggregating and averaging the local average output across workers per class. The resultant global average outputs are downloaded by each worker. Finally, to transfer the downloaded global knowledge into local models, each worker runs local iterations with its own loss function in addition to a regularizer measuring the gap between its own prediction output of a training sample and the global average output for the given class of the sample. Such a regularization method is called knowledge distillation (KD) that is to be elaborated in Sec. 5.2.

In contrast to FL and GADMM postulating the same model architecture for all workers, FD is capable of coping with heterogenous models. In fact, a widely-known application of KD is model compression, through which the knowledge of a large pre-trained model is transferred to an empty small model [67]. This is viable thanks to comparing the outputs of two models, regardless of their model architectures. FD inherits the same principle from KD, thereby allowing the workers with heterogeneous model architectures to be collectively trained. Compared to KD, FD is beneficial for preserving data privacy. In KD, the output of a pre-trained model may leak the raw data used during the pre-training process [68]. For instance, a generative adversarial network (GAN) can be trained under the guidance of the pre-trained model's outputs, such that the generator of GAN can reproduce synthetic samples that resemble the raw data [69]. One way to ameliorate this problem is to exchange ensemble of outputs, which leaks less information on raw data than individual outputs [70]. In FD, the outputs of each model are doubly averaged, i.e., locally within class and globally across workers, preserving more data privacy. For more details on the fundamentals of FD, readers are encouraged to check [71].

The effectiveness of FD is not limited to simple classification tasks under a perfectly controlled environment. In [72], FD is extended to an reinforcement learning (RL) application by replacing the aforementioned pre-class averaging step of FD with an averaging operations across neighboring states for an RL task. In [20], [73], [74], FD is implemented in a wireless fading channel, demonstrating comparable accuracy under channel fluctuations and outages with much less payload sizes compared to FL. Nonetheless, FD is more vulnerable to the problem of non-IID data distributions. Even if a worker obtains the global average outputs for all classes, when the worker lacks the samples in a specific target class, the global knowledge is rarely transferred into the local model of the worker.

3.4 Split Learning (SL)

A large-sized deep NN cannot be fit into edge devices' small memory. Split learning (SL) resolves this problem by dividing a single NN into multiple segments and distributing the lower segments across multiple workers storing raw data [21], [75]. By connecting the lower segments with a shared upper segment stored at a parameter server, each device uploads its NN activations of the cut-layer (i.e., lower segment's last layer) to the server calculating the loss values,

and downloads the gradients to update its lower segment. As done in FL, FD, and GADMM, SL does not exchange raw data, thereby partly preserving data privacy. To further reduce any possible private data leakage, differentially private mechanisms can be additionally applied before exchanging activations and gradients, e.g., by injecting fine-crafted noise [76]–[78]. With these benefits, SL has recently been adopted in medical applications wherein dispersed private health records should be exploited without sharing raw data [21], [75] [79]. SL has also been known for its robustness against non-IID data distributions, and applied for fusing heterogeneous vision and radio-frequency (RF) modalities to predict millimeter-wave channels [80]–[82].

While effective in terms of accuracy, the communication efficiency of SL is still questionable. As opposed to FL, FD, and GADMM that periodically exchange model updates, SL requires to exchange instantaneous model updates in feed-forward and backward propagations. For some applications, SL yields less communication overhead compared to the aforementioned periodic-update benchmark schemes by achieving much faster convergence [83], which may not always be feasible under different tasks and datasets. Furthermore, the communication cost of SL depends on the NN architecture and how to cut its NN layers, calling for more investigation on co-designing its communication and NN architectures.

3.5 Multi-Agent Reinforcement Learning (MARL)

Thus far we have implicitly considered that the learning objectives are focused on regression or classification tasks in supervised/unsupervised manner. When the environment dynamics effect on workers’ decisions, workers require to learn such dynamics and adapt their strategies through the experience obtained from worker-to-environment and inter-worker interactions. Towards this, reinforcement learning (RL) with the capabilities of exploration and exploitation plays a pivotal role. In RL, exploring allows agents to learn the dependencies of their choices on the environment and/or other agents (policy) as well as on the outcomes (value), which are then exploited to enhance the long-term rewards. Even in single agent scenarios, the required data to learn policy and value could be distributed over different workers that act as helpers. Under the limitations of communication and privacy, learning policy and value over distributed helpers could be benefited by FL, FD, and GADMM.

Interactions of multiple workers in a common environment while making decisions based on local observations are analyzed within the paradigm of multi-agent RL (MARL). Depending on the existence of a central controller and the types of interactions, MARL is categorized into centralized/decentralized and cooperative/competitive frameworks, respectively [84]. Centralized MARL frameworks postulate a central controller that learns decision-making policies by collecting all workers’ experiences that comprises their observed states, taken actions, and received rewards [85]. Exchanging such information may incur huge communication and memory resources while violating data privacy. Decentralized MARL without the central controller does not incur such issues, at the cost of not guaranteeing the equilibrium of the constituted policies of individual workers.

Even under cooperative MARL wherein all workers aim to achieve the same goal, it may not guarantee the convergence to equilibrium policies without central coordination [86]. Competitive MARL aggravates the problem, wherein every worker’s goal competes over a shared common environment and resources as a zero-sum game. Guaranteeing the convergence should thus require additional communication, as we shall discuss with a use case in Sec. 6.7. Nonetheless, note that all the rest of the discussions in this work are centered around distributed learning scenarios that are cooperative and NN based, rather than exploiting MARL in depth.

4 KEY COMMUNICATION PRINCIPLES

Both communication efficiency and accuracy of distributed learning can be improved by leveraging advanced communication principles coping with limited resources and time-varying communication dynamics as discussed in Sec. 2. Towards improving vanilla distributed learning methods presented in Sec. 3, several key communication principles are introduced in this section, and their effectiveness will be elaborated with selected use cases in Sec. 6.

4.1 Sparsification

Reducing the number of links can significantly decrease the communication bandwidth and energy of distributed learning. Such link sparsification can be implemented in temporal and/or spatial domain. Lazy aggregated gradient descent (LAG) [87] is one of its kind pursuing temporal link sparsity by enforcing each worker not to share its model update if the difference, measured by the infinity norm, between the current and previous updates does not exceed a certain threshold. Alternatively, to achieve the spatial link sparsity, one can enforce a sparse network topology by making each worker communicate only with very few neighbors, as exemplified by decentralized gradient descent (GD), dual averaging [88], and GADMM algorithms [65].

Link sparsification is not always free, but may come at the cost of higher learning convergence speed and/or lower accuracy. To illustrate, for the spatial link sparsification, a very sparse network graph (e.g., ring topology with nearest-neighbor based connectivity) yields high communication efficiency per iteration, but may incur more iterations for reaching the convergence and/or a target accuracy level, compared to a denser network graph (e.g., fully connected or star topology with the parameter server). Optimizing the sparsity under the trade-off between per-iteration communication cost and convergence speed is thus crucial.

While link sparsification reduces the bandwidth and energy, the per-link communication overhead still remains the same, which can be decreased by sparsifying gradient or model vectors. To be precise, model/gradient sparsification enforces transmitting n out of d elements at iteration k , where d is the total number of element in the gradient/-model vector and n is the number of the most important elements to send at iteration k . Gradient sparsification with fixed degree of sparsity approach were analyzed in [89]. On the other hand motivated by the fact that optimal sparsity depends on characteristics of the FL task, the computation

resources, and the communication bandwidth, [90] proposed an adaptive gradient sparsification scheme that can choose the n transmitted elements such that the training time is minimized.

4.2 Quantization

For each communication round, quantization decreases the number of bits to represent model updates, thereby reducing the communication payload sizes in distributed learning. Due to the reduced arithmetic precision of the model updates, quantization introduces errors, which may hinder the convergence of learning algorithms and/or degrade accuracy. Therefore, a quantizer and its quantizing levels should be carefully designed so as to guarantee the convergence with high accuracy. To this end, one can quantize each element of a gradient vector [91]–[93] or the gradient difference vector between the current and previous model updates [87], [94]. For the gradient quantization, the methods in [91], [92] adjust the quantizing levels under the trade-off between per-iteration communication cost and the convergence speed. SignSGD [93] considers an extreme case wherein gradients are quantized using only $+1$ and -1 , and shows its convergence by the aid of a majority vote of the workers. There are many other variants of gradient quantized distributed learning algorithms including error compensation [95], variance-reduced quantization [96], and ternary quantization [97].

Quantization can create synergy by integrating with link sparsification elaborated in Sec. 4.1. Lazily aggregated quantized gradient method (LAQ) is one example that combines the gradient update quantization with temporal sparsification, in a way that the number of links is sparsified based on the temporal gradient update difference, and the gradient update difference is adaptively adjusted for reducing per-link payload size while ensuring the convergence [87]. On the other hand, the method in [98] merges stochastic quantization with the spatial sparsification of GADMM [65], in which the weight update difference is rounded up and down with probability p and $1 - p$, respectively, while p is adaptively adjusted to minimize communication cost while preserving the convergence guarantees of vanilla GADMM [65].

The aforementioned methods quantize each element of a model update individually. Alternatively, the model update vector can be quantized altogether by clustering and mapping the updates into the centroids in a multi-dimensional vector space. Leveraging the universal quantization algorithm [99], the work [100] applies universal vector quantization to federated learning, coined UVeQFed, such that the quantization error can be bounded by a term that vanishes as the number of worker grows. Note that quantization can also be integrated with model/gradient sparsification in Sec. 4.1, under which the convergence guarantees are studied in [101].

4.3 Short Packet Aggregation

Whether the length of a communication packet is long or short has a significant impact on communication data rates. To be specific, in a large packet regime, the data rate R can be formalized by the well-known Shannon formula

$R = \log(1 + \text{SNR})$ per unit bandwidth over the additive white Gaussian noise (AWGN) channel for a given signal-to-noise ratio (SNR). Its derivation relies on assuming an infinite packet length $n \rightarrow \infty$ to ensure the zero packet error probability $\epsilon \rightarrow 0$ [102], and thus becomes a tight approximation for large packets. Since packet lengths are proportional to communication payload sizes, in the distributed learning context, the Shannon formula is suitable for deep NNs with periodic model parameter exchanging methods such as FL.

By contrast, SL exchanges a single NN layer’s instantaneous activation and gradient whose corresponding packet length can be very short. In this short packet regime with finite n and non-negligible ϵ , the data rate $R(n, \epsilon)$ can be described using a formula proposed by Y. Polyanskiy et al. [103], given as:

$$R(n, \epsilon) = \log(1 + \text{SNR}) - \sqrt{\frac{V}{n}} Q^{-1}(\epsilon) + \mathcal{O}\left(\frac{\log n}{n}\right), \quad (1)$$

where $Q(\cdot)^{-1}$ is the inverse of the Gaussian Q function, and V is the term capturing channel dispersion, e.g., under the AWGN, $V = (2\text{SNR} + \text{SNR}^2)(\log e)^2 / (1 + \text{SNR})^2$. This formula implies that the short packet length n incurs a penalty on the data rate that is proportional to $1/\sqrt{n}$. To alleviate such a penalty, one can aggregate consecutive packets, increasing n [104]. Through the lens of SL, this packet aggregation coincides with increasing the batch size of each worker. A larger batch size often yields faster convergence at the cost of compromising accuracy [105]. Consequently, there exists a trade-off among data rate, batch size, and accuracy in SL, as we shall discuss in Sec. 6.11.

4.4 Uncoded Transmission

The limited communication bandwidth is one key challenge in distributed learning over wireless channels. The wirelessly connected workers using the same channel may interfere with each other during their over-the-air transmissions. To avoid their interference, under digital transmissions, it is common to avoid such interference by allocate orthogonal channel bandwidths to different workers [13], [106]–[109]. As a result, the workers compete over the limited bandwidth, which is thus not scalable for supporting a large number of workers. Alternatively, motivated by the fact that the parameter server in FL is interested in the aggregated model updates of all workers, i.e., global model $\Theta = \frac{1}{N} \sum_{n=1}^N \theta_n$ with N workers, rather than the individual updates θ_n , several recent works have utilized uncoded transmissions, so as to harness interference without separate channel allocation [63], [110]–[115].

The idea of uncoded transmissions dates back to analog modulations and transmissions before the introduction of digital communications that are widely used in modern communication systems [22]. The current off-the-shelf systems by default do not support analog modulations. For analog communications, one may thus install additional analog modulators, and add a matched filtering functionality to decode the received analog signals under noisy channels [63], [110], [114]. Another way is to treat a digital modulator as an analog modulator integrated with quantization [116]. More feasibility of supporting uncoded

transmissions against fading, synchronization, and many other issues is discussed in [117]. By nature, uncoded transmissions are prone to signal distortions and perturbations. Nonetheless, in the context of ML, noise is often useful in regularization [118], which can therefore be partly tolerable without compromising accuracy [116]. For excessive noise, power control and sophisticated optimization methods can mitigate possible accuracy degradation as elaborated next.

Under analog transmissions, each transmitted signal from a worker in FL is perturbed by fading, i.e., multiplied by the fading gain h_n , and superpositioned over-the-air with all other workers' signals using the same channel. Consequently, $\frac{1}{N} \sum_{n=1}^N h_n \theta_n$ is received by the parameter server. The superpositioning property of analog transmissions is favorable for averaging the models updates using the entire bandwidth for all workers, rather than competing over the limited bandwidth with each other under digital transmissions. By contrast, the fading perturbation may hinder obtaining the received signal in a desired form at the parameter server, e.g., equal or weighted averaging with the weight that is proportional to the ratio of each worker's data size [119]. One way to cope with the fading perturbation is the channel inversion method. By inversely perturbing the signal before transmission, i.e., multiplying by $1/h_n$, the fading can be canceled out at reception [114]. This channel inversion however consumes the transmit power inversely proportional to the channel gain, which is not viable for small h_n under the limited edge device energy budget. For this reason, it is common to allow transmissions only when the channel gains exceed a certain threshold [110]–[112]. As discussed in Sec. 4.1, such temporal sparsification may hinder the convergence of learning algorithms.

Alternatively, the method proposed in [63] only utilizes the superpositioning property of analog transmissions without channel inversion. This is done by reformulating FL and optimizing it directly with perturbed model updates as follows. To be specific, recall the original unconstrained problem of FL, aiming to minimize $\frac{1}{N} \sum_{n=1}^N f_n(\Theta)$, by locally minimizing $f_n(\theta_n)$ at each worker and globally averaging their model parameters θ_n at the parameter server, yielding Θ . This boils down to the following constrained average consensus problem:

$$\min_{\Theta, \{\theta_n\}_{n=1}^N} \sum_{n=1}^N f_n(\theta_n) \quad (2)$$

$$\text{s.t. } \theta_n = \Theta, \quad \forall n. \quad (3)$$

To incorporate the fading perturbed model updates in the problem formulation, by multiplying the fading gain h_n at both sides, (3) is recast as its equivalent constraint $h_n \theta_n = h_n \Theta, \forall n$. This reformulated problem is solved using ADMM while directly incorporating the perturbed model updates, i.e., $h_n \theta_n$, without inverting the fading gain h_n . As a consequence of avoiding channel inversion, the convergence becomes less sensitive to the transmit power constraint. Furthermore, thanks to directly exploiting the perturbed model updates, it is more robust against the adversarial or honest-but-curious parameter server, to be further elaborated in Sec. 6.4.

Still, analog multipath, channel dispersion, synchronization errors

4.5 Scheduling and Offloading

Heterogeneity is prevalent in distributed learning, in terms of the availability and access to the training data and resources for the communication, computation, and memory. Such heterogeneity results in the learning workers having outdated models compared to other workers, referred to as *stragglers*. Waiting these stragglers may cause significant delays to the overall training operations, whereas ignoring them may hinder guaranteeing the convergence or achieving high accuracy. Scheduling is effective in balancing and resolving this straggler handling problem. To this end, it is of paramount importance to identify the root cause of each straggler and its contribution to the overall learning performance.

The lack of computing resources can be one major cause of stragglers. It happens when large-sized models and datasets with multiple tasks are processed by on-device and battery-limited workers. In this case, as studied in [120], an effective solution could be scheduling the resultant stragglers while offloading their computationally demanding tasks (or even training data with a loss of privacy) to neighbors or edge servers, a conceptual design known as mobile edge computing (MEC) [121], [122]. Such task offloading in MEC needs to take into the account of device heterogeneity [123], communication limitations [124], [125], and demand-supply capabilities of processing power [126] in addition to its impact on the tolerable training latency [122] and target training/inference accuracy [127] while ensuring devices' privacy [128].

Another source of stragglers is poor channel conditions such as the channels in deep fades and severe interference, as well as communication resource limitation such as limited bandwidth and uplink transmit power. Especially, frequent and simultaneous computation and communication are impractical within the large-scale systems, in which resources management via client scheduling is essential for the identification, coordination, and diminution of the aforementioned stragglers [129]–[134]. In addition to client scheduling and resource management, to remedy this type of straggler problem, it is useful to utilize advanced multiple access control techniques such as proactive scheduling via channel prediction and multi-hop relaying [72], [135], [136].

5 KEY MACHINE LEARNING PRINCIPLES

Communication efficiency of distributed learning is significantly affected by ML architectures and algorithms. In this section, several machine learning principles are presented for improving vanilla distributed learning methods discussed in Sec. 3, and their effectiveness will be validated by representative use cases in Sec. 6.

5.1 Model Split

Running a large-sized deep NN consumes huge memory that may not fit within edge devices. The energy consumption of this model is proportional to the model sizes [137], aggravating the problem under battery-limited edge devices. SL resolves such issues by splitting a single NN model into multiple segments stored and operated by different edge nodes. In essence, this problem is traced back to

model parallelism focusing on how to partition and offload NN segments, as opposed to data parallelism considering a large-sized global dataset dispersed across different workers running NN models, each of which is separate but has the same architecture [138].

Traditionally model parallelism has focused primarily on the NN partitioning based on computing latency [139]. For instance, a convolutional NN comprises fully-connected layers and convolutional layers, and the convolutional layers often consume much longer processing delays compared to the fully-connected layers, e.g., in AlexNet [140] and ResNet [141] architectures. Therefore, even if two edge nodes have the same memory size, equally partitioning an NN may not be an optimal way, incurring imbalanced processing overhead. Beyond this, in the context of SL, communication efficiency and data privacy should also be taken into account. Indeed, cutting a NN’s bottleneck layer having the smallest dimension (e.g., VAE’s bottleneck layer for latent variables [142]) can maximally reduce the SL communication payload sizes. Furthermore, in a classification task, not only unlabeled data samples but also their ground-truth labels can be privacy-sensitive (e.g., unlabeled X-ray images and their ground-truth diagnosis results) [143]–[145]. In this case, the input and output layers are linked to the raw samples and ground-truth labels (for training loss calculation), respectively, and a NN should thus be partitioned such that the input and output layers are colocated at the data owner. More discussions on model split are deferred to Sec. 6.10.

5.2 Knowledge Distillation

Knowledge distillation (KD) aims to imbue an empty student model with a teacher’s knowledge [67]. In a classification task, as opposed to the standard model training that attempts to match the student model’s one-hot prediction (e.g., [cat, dog] = [0,1]) of each unlabeled sample with its ground-truth label, KD tries to match the model’s output layer activation, so-called logit (e.g., [cat, dog] = [0.3, 0.7]), with the teacher’s logit for the same sample. This logit contains more information than its one-hot prediction, thereby training the student model faster than the standard training with much less samples [146].

The teacher’s knowledge of KD can be constructed in different ways. Originally, the knowledge is a pretrained teacher model’s logit, which is transferred to a smaller student model for model compression [67]. The knowledge can also be an ensemble of other student models’ logits [19], in that the ensemble of predictions is often more accurate than individual predictions. Leveraging this to enable distributed learning, the knowledge in FD is constructed the ensemble of different workers’ prediction, each of which is locally averaged per label in a classification task [143], [147] or across neighboring states in reinforcement learning [148]. The local averaging step avoids the same sample observations of the student and teacher models (i.e., ensemble of all student models), thereby reducing significant communication overhead while preserving local data sample privacy. Lastly, for given averaged logits as the teacher’s knowledge, running KD with an empty student model at the parameter server realizes a fast one-shot FL or the information type conversion from logits to the parameters of the trained

student model, which will be discussed with a use case in Sec. 6.11.

5.3 Mixup Augmentation

Mixup is a data augmentation technique generating a synthetic sample by superpositioning two different samples [36]. As an example, in a binary classification task, a sample s_0 in the label 0 is linearly combined with another sample s_1 in the label 1, thereby yielding a synthetic sample \hat{s}_{01} given as:

$$\hat{s}_{01} = \lambda s_0 + (1 - \lambda) s_1. \quad (4)$$

The term λ is the mixing ratio that is randomly sampled from a bathtub-shaped beta distribution such that \hat{s}_{01} resembles a sample in the label either 0 or 1 with a slight difference. Manifold Mixup applies the same technique to superposition two different hidden representations, which often performs similar or even higher accuracy than vanilla Mixup that combines raw samples [149].

Both vanilla and Manifold Mixup are commonly used in standalone training, particularly for adversarial learning that intentionally feeds distorted samples to obtain more generalized models [36], [149]. In distributed learning, these techniques can also be utilized for sharing proxy samples without revealing raw data samples [150]. For example, to rectify non-IID data distributions, each worker can exchange mixed-up samples or manifold mixed-up representations to complement missing samples in some labels [144]. By uploading the mixed-up samples or representations to a parameter server, the workers’ training computation can be offloaded to the server enabling one-shot FL [151]. The number of these generated proxy samples or representations can further be oversampled by mixing them across different workers [148] and/or re-mixing the mixed-up samples or representations [73]. More use cases and effectiveness of Mixup and manifold Mixup will be discussed in Sec. 6.8, 6.9, and 6.12.

5.4 Gaussian Process Regression

Dynamics in the environment, agents’ hardware, and random choices of training batches and learning parameters cause computing and communication resources and training model parameters to change over the training duration. These dynamics of resource and model states can be viewed as stochastic processes. Considering a Gaussian process prior probability distribution on above stochastic processes provides means of analyzing them using Bayesian inference methods [152]. Gaussian process regression (GPR) is the process of determining a set of kernel hyperparameters defining the covariance matrix between the all possible observations over time (and space) assuming zero-mean distribution therein. Using GPR, the posterior mean and variance at unseen observations can be analytically estimated.

By modeling the dynamics of the resource (computation and/or communication) availability as a time series, GPR can be adopted to predict future resources (mean) with the uncertainty bounds (variance) [153]. This allows to identify agents who are likely to be stragglers in advance, in which agents and resources can be proactively scheduled. As a result, the overall training latency can be decreased

with minimum loss of training performance and the overall resource utilization can be improved. Similarly, model parameter dynamics can be analyzed with GPR. Under the communication bottleneck in collaborative learning, agents can utilize estimated model parameters of others to continue local training while using the limited resources only when the uncertainties of model estimations are unacceptable.

5.5 Mean-Field Game (MFG) Learning

Decentralized decision-making of competitive and mutually interactive workers is a challenging task as discussed in Sec. 3.5. Due to these interactions, it is common to determine a single worker's action by fixing all other worker states, and then iterate it for the next worker until all workers' actions converge to the Nash equilibrium, a stable state at which no worker gains more reward by changing its action [154]. The complexity of this problem is thus increasing exponentially with the number of workers, which is unfit for dealing with massive interactive workers. Mean-field game (MFG) is a useful framework to greatly reduce the complexity [155]–[159]. At its core, MFG approximates the problem of massive interactive workers as the problem of each single worker interacting with a virtual worker whose state distribution is given by the distribution of the entire population. Then each worker's decision-making boils down to solving two partial differential equations (PDEs), the Hamilton-Jacobi-Bellman (HJB) equation and the Fokker-Plank-Kolmogorov (FPK) equation [160]. By solving FPK, one can obtain the population state distribution, called mean-field (MF) distribution. For the given MF distribution, solving HJB results in the optimal action of each worker.

One common limitation of MFG-theoretic approaches is the curse of dimensionality, which is detoured by the MFG learning framework. To be specific, a PDE is often solved numerically by discretizing the domain so that the derivatives therein can be approximated using finite differences. To guarantee the convergence of such a finite difference method, the discretizing step size should decrease with the domain dimension. As an example, for a given n -dimensional domain vector $\{x_1, x_2, \dots, x_n\}$, the discretization step size Δt should satisfy $\Delta t \leq (\sum_{i=1}^n 1/x_i)^{-1}$ according to the Courant-Friedrichs-Lewy condition [161]. Consequently, the dimensionality increase in states and actions incurs huge extra computing overhead for solving FPK and HJB equations, respectively. MFG learning resolves this issue by recasting the problem of solving HJB and FPK equations, i.e., $H = 0$ and $F = 0$, respectively, as the regression tasks of minimizing $|H|$ and $|F|$, respectively. To solve these two regression tasks, a pair of HJB NN and FPK NN are introduced in that NNs are good at tackling regression problems via simple first-order algorithms such as the gradient descent method. The effectiveness of MFG learning will be corroborated with a massive drone control use case in Sec. 6.7.

5.6 Convergence of FL

Asymptotic convergence analysis on a distributed learning algorithm sheds a light on its design guidelines. For instance, it can provide an insight on how to accelerate the training process by most efficiently exploiting the limited

computing and communication resources under a given network topology, channel dynamics, and data distributions. To this end, the fundamentals of FL have recently been studied through the lens of the convergence analysis on local SGD [162]–[165]. Compared to the standard minibatch SGD averaging each model per iteration [166]–[168], local SGD locally runs multiple iterations, followed by averaging the models across workers, which is identical to FedAvg [169]. With M workers during T iterations, it has been shown by [163] that local SGD can achieve the convergence rate of $\mathcal{O}(1/\sqrt{MT})$ under both IID and non-IID data. This implies a linear speed up in M , i.e., increasing M is equivalent to increasing T . To achieve such a linear speedup, the result clarifies that under non-IID data one should communicate more frequently with the communication interval $H = 1 + \lfloor T^{1/4} M^{-3/4} \rfloor$, compared to $H = \mathcal{O}(T^{1/2} M^{-3/2})$ under IID data for $T > M^3$. Similar results have also been shown in [164], [165], proving that under IID data, local SGD improves over minibatch SGD in terms of the worst-case errors, whereas under non-IID data the opposite is true in most cases.

With the aid of a central parameter server and ideal communication channels without resource limitations, distributed learning methods including FL commonly guarantee fast convergence. Indeed, minibatch SGD achieves the linear convergence rate $\mathcal{O}(\log(1/\varepsilon))$ for the optimality gap ε , even with quantization [170]. By contrast, without the aforementioned ideal assumptions, convergence guarantees become no longer trivial. For a decentralized architecture under a sufficiently connected network topology, a sophisticated primal-dual algorithm has been proposed, which achieves the linear convergence rate with quantization [171]. For a sparse bipartite graph including ring network topologies, GADMM and its variants have been proposed for achieving the linear convergence under quantization, link censoring, and time-varying topologies [65], [98], [172]. Another roadblock to the convergence is limited communication resources and poor channel conditions particularly under wireless connectivity. To overcome long-distance path loss, a multi-hop FL method has been proposed, and its convergence has been studied in [29]. Under fading channels with limited bandwidth, the convergence bound has been derived, through which the communication resource allocation has been optimized in [173], [174].

5.7 Convergence of FD

The convergence properties of KD and FD have recently been investigated in [71], [175] by exploiting the theory of neural tangent kernel (NTK) [176]. NTK is based on the empirical finding that a large-width model and its random initialization jointly induce a kernel regime under which the model output is represented in a closed-form expression. This kernel regime enables an asymptotic analysis on KD that compares a pre-trained teacher's and student's outputs as elaborated in Sec. 5.2. As the number of iterations T goes to infinity, the gap between the ground-truth \mathbf{y} and the student's output $\mathbf{f}(T)$ in KD is thereby given by [175] as:

$$\left\| \lim_{T \rightarrow \infty} \mathbf{f}(T) - \mathbf{y} \right\|_2 = A \left\| \mathbf{y} - \underbrace{\sum_{m=1}^M \frac{a_m \phi_m}{\sqrt{M}}}_{\text{teacher's output}} \right\|_2, \quad (5)$$

where M is the model's width, $A = \lambda/(a + \lambda)$ with $|A| < 1$, $a = \sum_{m=1}^M a_m/M$, and $\{\phi_m\}$ is the logit that is combined by the weight vector $\{a_m\}$ to construct the model's output. In (5), the RHS decreases as the teacher's output approaches to \mathbf{y} .

Towards extending the NTK analysis of KD in (5) to that of FD without any pre-trained model, for mathematical amenability, all models are supposed to synchronously observe the same data samples as assumed in [19]. This can be regarded as the upper bound performance of FD in that the assumption yields more accurate model output comparison than the original FD comparing each model's output with the globally averaged output per class. For N workers with such an assumption, as $T \rightarrow \infty$, the gap between the ground-truth \mathbf{y} and the output $\mathbf{f}^1(T)$ of a randomly selected worker in FD is derived in [71] as:

$$\left\| \lim_{T \rightarrow \infty} \mathbf{f}^1(T) - \mathbf{y} \right\|_2 = \frac{A}{\lambda} \left\| \alpha A^r + \beta \left(\frac{A}{N-1} \right)^r \right\|_2 \stackrel{r \rightarrow \infty}{\approx} 0, \quad (6)$$

where r is the number of communication rounds, and the finite terms α and β are specified in [71]. The asymptotic result implies that FD not only guarantees the convergence but also has a potential to outperform KD whose performance is upper-bounded by how well the teacher model is pre-trained as shown in (5).

For other key distributed learning algorithms discussed in this article, the convergence properties of GADMM and its variants will be elaborated in Sec. 6.1, 6.2, and 6.3. Under ideal channels, the training dynamics of vanilla SL is identical to standard minibatch SGD [166], [167], and the convergence properties also follow those of minibatch SGD [168]. Still, there are many missing convergence analyses such as SL under outage channels and FD under finite-width model architectures and decentralized topologies, which could be interesting topics for future research.

6 USE CASES: COMMUNICATION-EFFICIENT AND DISTRIBUTED LEARNING FRAMEWORKS

By applying the ML and communication principles introduced in Sec. 4 and 5 to vanilla distributed ML methods in Sec. 3, in this section we present communication-efficient and distributed learning frameworks with selected use cases. The mapping between specific principles and use cases is illustrated in Fig. 1.

6.1 Quantized-GADMM (Q-GADMM)

Utilizing GADMM that exploits sparse connectivity (Sec. 4.1), Q-GADMM allows each worker to share a quantized version of its model with neighbors [98]. Using stochastic quantization, one of the key communication principle in Sec. 4.2, with adjustable quantization range, Q-GADMM can significantly reduce the communication energy compared to original GADMM at a zero cost in terms of the convergence speed and accuracy.

The stochastic quantization places the i -th dimensional element $[\hat{\theta}_n^{k-1}]_i$ of the previously quantized model vector at the center of the quantization range $2R_n^k$ that is equally divided into $2^b - 1$ quantization levels. This yields a quantization step size $\Delta_n^k = 2R_n^k/(2^b - 1)$ of resolution b . Each




Fig. 2. Q-GADMM: relation between energy consumption and relative linear regression loss ($|F - F^*|$).


worker quantizes the difference between the current and the previously quantized models by choosing a rounding probability yielding a zero quantization error on average.

Each worker then transmits R_n^k and the index of the quantization level $q_n(\theta_n^k)$ to its neighboring workers. At the receiver, $\hat{\theta}_n^k$ can be reconstructed by $\hat{\theta}_n^k = \hat{\theta}_n^{k-1} + \Delta_n^k q_n(\theta_n^k) - R_n^k \mathbf{1}$. Consequently, when the full arithmetic precision uses 32 bits to represent R_n^k , the payload size of Q-GADMM is $(bd + 32)$ bits where d is the model size. Compared to GADMM whose payload size is $32d$ bits, Q-GADMM can achieve a huge reduction in communication overhead, particularly for large d .


Fig. 2 compares Q-GADMM with GADMM, and two PS-based schemes (QGD, and ADIANA [177]) in terms of the loss versus the total sum energy for a system of 50 workers. Here, linear regression of California housing dataset with $d = 6$ input features is tested. In the full precision GADMM, each worker will transmit $32d$ bits to represent all elements in the model vector. In contrast, each worker of Q-GADMM only uses $(32 + 2d)$ bits, with 2 bits to represent each element in the model vector. Following the Shannon's capacity theorem, more bits consumes more transmission energy for the same bandwidth, transmission duration, and noise spectral density. Fig. 2 exhibits significant reduction in the total energy consumption, a key challenge discussed under Sec. 2, for Q-GADMM compared to all baselines, owing to i) the decentralization where workers communicate with only nearby neighbors (Sec. 4.1), ii) the fast convergence inherited from GADMM (Sec. 3.2), and iii) the reduction of transmitted bits at every iteration while ensuring convergence via stochastic quantization (Sec. 4.2).

6.2 Dynamic GADMM (D-GADMM)

In practise, due to device mobility, the network topology is time variant, in which neighboring nodes continuously change over time as highlighted in Sec. 2. Hence, to enable distributed learning over dynamic network of workers, Dynamic GADMM (D-GADMM), which inherits the theoretical convergence guarantees of GADMM is proposed in [65]. While adapting to network dynamics, D-GADMM improves the convergence speed of GADMM, i.e., random changes in



(a) Communication iterations.



(b) Communication cost.

Fig. 3. D-GADMM: comparison of objective errors as functions of (a) communication iterations and (b) communication cost.

sparse and logical neighbors (Sec. 4.1) of a static physical topology can significantly accelerate the convergence of GADMM. Although the sparsity of network graphs yields slow convergence speeds [178], the reductions of convergence speed in D-GADMM compared to the standard PS-based ADMM can be compensated by continuously altering neighbors with D-GADMM. In addition, with dynamic topology changes, D-GADMM exhibits significant communication cost reductions compared to GADMM [65].

Fig. 3 compares D-GADMM with both GADMM and standard ADMM. From Fig.3, it can be seen that utilizing D-GADMM significantly increases the convergence speed of GADMM and hence, reduces the total communication cost even when the topology is fixed. Therefore, D-GADMM can compensate for the decrease in the convergence speed of GADMM compared to PS-based ADMM due to topology decentralization and maintains a low communication cost per iteration gained by GADMM.

6.3 Censored Generalized GADMM (C-GGADMM)

In GADMM, every worker has to share its own model with only up to two neighboring workers at every iteration. To reduce communication overhead while addressing more general network topologies, we propose censored generalized GADMM (C-GGADMM). In C-GGADMM, by





Fig. 4. C-GADMM: loss as a function of total energy consumption.

exploiting temporal sparsity, each worker shares its model only if the difference between the current and the previous models exceeds a certain threshold [87]. Furthermore, each worker in C-GGADMM can communicate with an arbitrary number of neighbors in a different group (i.e., under any bipartite graph), which is helpful addressing time-varying network topologies (Sec. 2). Theoretically, C-GGADMM inherits the same performance and convergence guarantees of Vanilla GGADMM, under a non-increasing and non-negative censoring threshold sequence; particularly if the threshold at iteration k follows $\tau_k = \omega\zeta^k$ where $\omega \geq 1$ and $0 < \zeta^k < 1$. Furthermore, by integrating Q-GADMM and C-GGADMM, we propose C-QGGADMM that performs the censoring based link sparsification with payload quantization (Sec. 4.2). Consequently, C-QGGADMM decreases both the cost per channel use and the number of channels, thereby significantly reducing the communication energy and competition on the limited bandwidth.

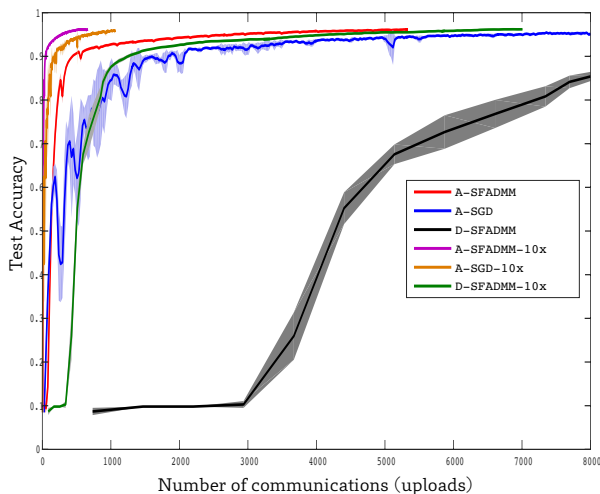
The benefits of censoring and quantization in terms of reduced energy consumption are elaborated in Fig.4 using the linear regression problem described in Sec. 6.1. It can be noted that introducing censoring on top of GGADMM can provide about two-fold reduction in the total communication cost. Moreover, implementing both censoring and quantization can further lower the total communication cost.

6.4 Analog Federated ADMM (A-FADMM)

In A-FADMM [63], each worker transmits an analog signal (Sec. 4.4) that is a function of the i -th element in the model over a *shared channel among all workers*. All transmitted signals are superpositioned over the air while hiding each private local model in the crowd preserving the privacy (Sec. 2). Consequently, the PS receives aggregated signals of all individuals *perturbed* by their complex fading channels. Hence, A-FADMM aggregates multiple workers' updates at the PS without competition on the available bandwidth via analog transmissions. It was proven in [63] that A-FADMM converges to the optimal solution for convex functions and preserves privacy. Moreover, A-FADMM copes with the nuisances incurred by analog transmissions, in terms of



(a) Linear regression.



(b) Image classification with a deep NN.

Fig. 5. A-FADMM: performance comparison in (a) a linear regression task and (b) the MNIST classification task with a deep NN.

time-varying channel fading, noise, and transmit power limitation.

Fig. 5(a) compares analog and digital implementations of ADMM on a linear regression task. We plot the loss vs the number of uploads (communication rounds). As observed in Fig. 5(a), A-FADMM requires the lowest communication rounds to achieve a target loss 10^{-4} . Even with $10\times$ more subcarriers, D-FADMM fails to reach the same speed due to the orthogonal subcarrier allocation to each worker under limited bandwidth. However, if one aims to achieve very low loss below 10^{-4} , A-FADMM suffers from noisy reception, and D-FADMM may thus be a better choice, as long as very large bandwidth and/or long uploading time are available.

Fig. 5(b) validates the applicability of the stochastic version of A-FADMM (A-SFADMM) on the stochastic and non-convex problem of image processing using DNN. Note that the model size for the tested DNN architecture is several orders of magnitudes higher than the model size of the linear regression problem discussed above (For the simulation details see [63]). As observed from Fig. 5(b), A-

SFADMM significantly outperforms the digital implementation (D-SFADMM) in terms of the convergence speed while achieving the maximum accuracy. In fact, A-SFADMM outperforms $10\times$ -D-SFADMM which has $10\times$ larger bandwidth (i.e., $10\times$ more subcarriers).

6.5 Quantum Scheduler Aided FL

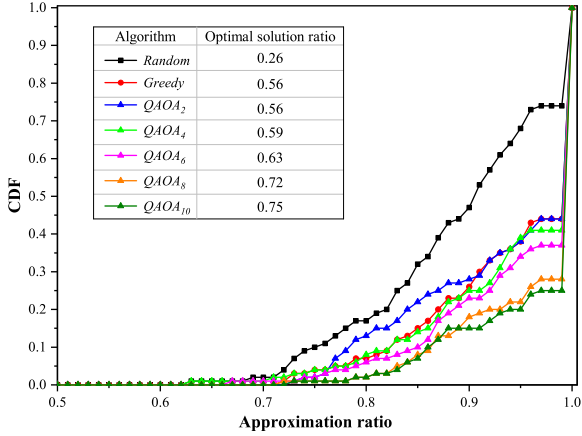
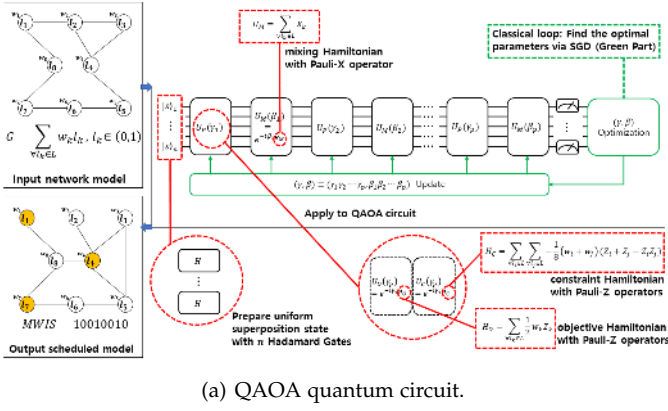
In modern quantum computing research, the design and implementation of quantum approximate optimization algorithms (QAOA) is of great interest [179], [180]. With QAOA-based methods, many approximation algorithms for NP-hard problems are under development. Among the NP-hard problems, the QAOA-based approximation solution approach to max-weight independent set (MWIS) problem is actively under discussion where the MWIS formulation is widely used for network scheduling modeling, e.g., device-to-device wireless networks [181]. As studied in [182], scheduling problems are considered and formulated with MWIS in FL over wireless channels where the objective for the scheduling is sum-rate-maximization. Compared to classical MWIS solution approaches such as message passing, QAOA is envisaged to be beneficial in terms of computation time and complexity, particularly with the advances in quantum computers.

To solve an MWIS problem, QAOA is described by a quantum circuit in [183] as shown in Fig. 6(a). The circuit aims to find proper parameters for quantum approximation of the MWIS formulation, which is a combinatorial optimization problem. The quantum circuit has p levels, in which increasing p yields more number of the alteration of quantum approximating computation, improving the accuracy at the cost of reducing the convergence speed. After finding the parameters, the solution to the MWIS problem is obtained from the optimum of the expected Hamiltonian. Such a QAOA quantum circuit is implemented in [183] by using Cirq and TensorFlow-Quantum [184], where Cirq is a Python framework for creating, editing, and invoking noisy intermediate scale quantum (NISQ) circuits, while TensorFlow-Quantum integrates quantum computing algorithm with the logic designed in Cirq.


As presented in Fig. 6(b), the QAOA-based MWIS scheduler (QAOS) outperforms greedy and random scheduling baselines, wherein the performance is measured by the cumulative distribution function (CDF) of the proportion between the scheduled workers' weights under QAOS and the optimal weights after an exhaustive search. The same tendency holds for various p values. Next, Fig. 6(b) shows the impact of QAOS on FL. When counting only non-interfering links out of 10 workers, FL under QAOS with $p = 2$ can serve 6.12 workers on average, while greedy and random schedulers support 6 and 2 workers on average, respectively. Consequently, in the MNIST classification task, FL with QAOS achieves the highest accuracy compared to other two benchmark schemes.

6.6 GPR Aided FL

Communication plays a key role in the local model aggregation and global model sharing steps of the FL (Sec. 3.1) over wireless networks. The poor channel conditions in both uplink and downlink introduces stragglers from the



(b) Performance of the QAOA based scheduler (QAOS).



(c) Test accuracy of FL with QAOS.

Fig. 6. QAOA-based quantum scheduler (QAOS) and its impact on FL: (a) a schematic illustration of the QAOA quantum circuit; (b) the CDF of the ratio between the QAOA-scheduled traffic of 10 workers and the traffic scheduled via an exhaustive search; (c) Test accuracy learning curve of FL with QAOS compared to FL with greedy and random schedulers.

communication point of the view (Sec. 2), in which, channel measurement or accurate estimation is essential under the limited communication resources [27]. Although measuring channels aids to utilize agent and resource scheduling, the




Fig. 7. Comparison of FL over MNIST dataset with GPR-based channel prediction and joint agent and wireless resource scheduling.

channel sampling and pilot transmissions therein require high reliable (possibly dedicated) resources as well as introduce significant latency to the training process. To overcome the cons of channel measurement, GPR-based channel estimation can be adopted in FL (GPR-FL) [130]. By modeling the dynamic channel states as stochastic processes with a Gaussian prior, time series prediction in GPR can be used to estimate the channels and their uncertainty (Sec. 5.4). Using the uncertainty of channels from GPR as a regularizer within FL loss function, joint channel sampling and allocation for straggler-free scheduling (Sec. 4.5) can be carried out simultaneously to reduce the sampling latency [185].

To illustrate the benefits of GPR-FL, we compare the training loss dynamics (relative to the loss of centralized training) of GPR-FL under limited wireless resources with three other methods as illustrated in Fig. 7: i) *SchedFL*: joint agent and resource scheduling towards minimizing training loss similar to GPR-FL is used with channel measurements, ii) *PF*: proportional fair scheduling in terms of contribution to model aggregation without channel measurements, and iii) *IDEAL*: FL without communication constraints. Note that a single resource block is dedicated for the channel measurement in SchedFL. GPR-FL reaps the benefit of the additional resource by utilizing it in agent scheduling over SchedFL, yielding a lower loss as close to IDEAL. In contrast, PF performs poorly even with the additional resource, due to the absence of the training loss minimization objective within its scheduling policy. It is worth noting that due to underlying complexity in GPR and lack of channel sampling, GPR-FL may tend to lose its performance under the availability of excessive amount of total resources compared to the number of agents. A viable solution is to limit agents' access to subsets of resources rather the entire resource pool.

6.7 Federated MFG Learning for Massive UAV Control

By integrating MFG learning with FL, in this use case we study controlling a massive number of unmanned aerial vehicles (UAVs) in a communication-efficient and decentralized way. Following the MFG learning framework as elaborated in Sec. 5.5, each UAV is equipped with a pair of HJB and FPK NNs. The HJB NN outputs (i) the UAV's optimal

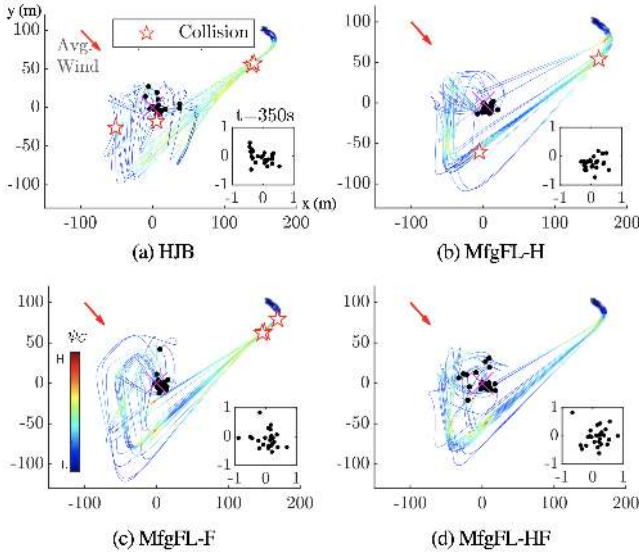


Fig. 8. Trajectories of 25 UAVs dispatched from a common source (top right) to a destination (bottom left), when each UAV runs MFG learning while exchanging: (b) only HJB NNs, (c) only FPK NNs, and (d) both HJB and FPK NNs with neighbors, compared to a baseline in which (a) each UAV runs only an HJB NN while exchanging raw states.

action (i.e., acceleration) and (ii) the resultant cost functional value by feeding (iii) the UAV's observed state and (iv) the state distribution of the entire UAV population (i.e., MF distribution). The FPK NN outputs (iv) the MF distribution by feeding (iii) the UAV's state and (ii) the cost functional value obtained from the HJB. While (iii) is fixed, (ii) and (iv) are recursively updated until convergence, at which the optimal action (i) is finally determined [158], [159]. According to the MFG theory [186], the aforementioned optimal control can achieve the epsilon-Nash equilibrium as long as the initial states of all UAVs are exchanged without any further inter-UAV communication. This is true when the outputs of the HJB and FPK NNs accurately approximate the solutions of the HJB and FPK equations; in other words, HJB and FPK NNs are ideally trained, which is not feasible due to the lack of training samples (i.e., observed states).

To accelerate the training of HJB and FPK NNs, following FL, each UAV periodically broadcasts its NN weights with its neighbors, and updates its model by averaging the received weights within a predefined latency deadline. As each UAV has HJB and FPK NNs, there are three possible configurations, exchanging only HJB NN (MfgFL-H), only FPK NN (MfgFL-F), or both HJB and FPK NNs (MfgFL-B) at the cost of the increased communication payload sizes. With 25 UAVs dispatched from a common source to a destination, Fig. 8 shows that MfgFL-B achieves the best trajectory without any collision, while all MfgFL based methods yield better results than a baseline operates by only running the HJB NN while exchanging raw states of neighboring UAVs. Here, the curve color indicates the value of ϕ_G in the cost function, a swarming term that decreases with the relative velocities, and increases with the relative distances of all UAVs. Again, MfgFL-B yields the lowest ϕ_G even at the early stage, supporting the collision-free results. Furthermore, MfgFL-B consumes the minimum motion energy until reaching the destination as shown in Fig. 9(a), and

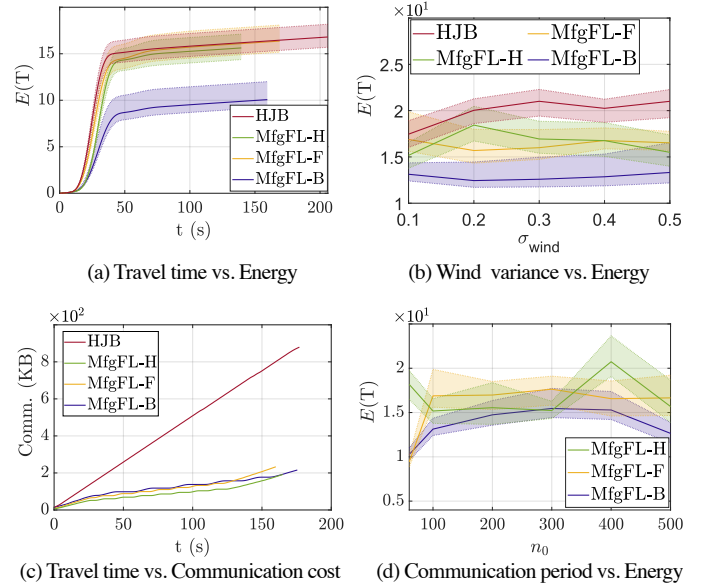


Fig. 9. Performance comparison of MfgFL in terms of: (a) energy and travel time, (b) energy for various wind velocity variance values, (c) accumulated communication cost during travel, and (d) energy for various communication periods.

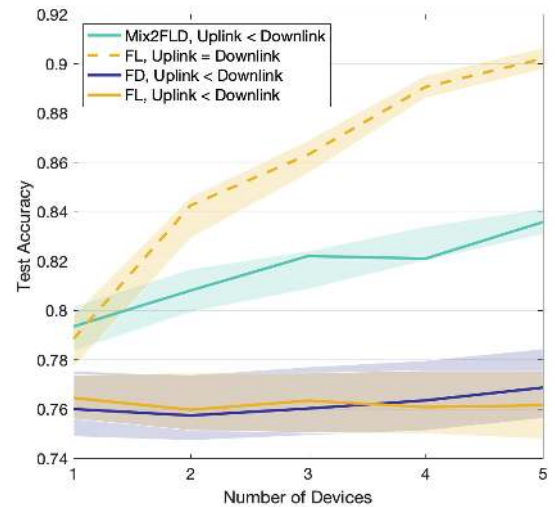


Fig. 10. Test accuracy of Mix2FLD for a different number of workers, compared with FL and FD under symmetric (Uplink=Downlink) and asymmetric (Uplink \neq Downlink) channel capacities.

is more robustness against external disturbances reflected by the variance of random wind velocity as observed in Fig. 9(b). In addition, Fig. 9(c) illustrates that MfgFL-B exchanges the least amount of the packets even though its per-communication payload size is 2x greater than MfgFL-H or MfgFL-F. Lastly, for different communication periods, MfgFL-B results in the least energy consumption as seen by Fig. 9(d).

6.8 Downlink FL After Uplink FD

In mobile communication systems, uplink data rates are often much lower than downlink rates due to the limited transmission power of mobile devices [187]. Therefore, FD

(Sec. 3.3) is useful in the uplink thanks to its small payload sizes, whereas in the downlink FL (Sec. 3.1) is preferable in that exchanging model parameters commonly achieves higher accuracy than exchanging model outputs [73]. To jointly exploit FD and FL under uplink-downlink asymmetric channels, we present an FL-after-FD algorithm combined with two-way Mixup (Mix2FLD). In Mix2FLD, the model outputs (i.e., logits) are uploaded to a server via FD, which should be converted into a global model whose parameters can be downloaded by and updated at each device using FL. Such a model output-to-parameter conversion is viable using KD (Sec. 5.2) that updates the global model at the server by minimizing the difference between the uploaded outputs and the the outputs of the global model. This requires a handful of seed samples to generate the global model’s outputs, which is a major challenge of its implementation due to the extra communication overhead and possible data privacy violation as highlighted in Sec. 2.

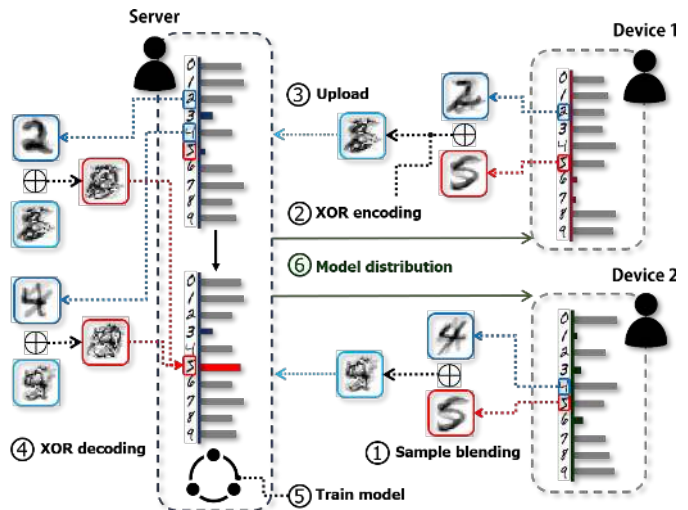
In a classification task, we resolve the aforementioned problem by applying the Mixup method twice. Precisely, before uploading each device encodes multiple samples by them via Mixup (Sec. 5.3). Then, the server decodes the Mixup-encoded samples uploaded from different devices by additionally superpositioning them, in a way that the decoded samples have one-hot labels. Such a decoding commonly improves accuracy particularly under non-IID data distributions [73], [151]. Note that the encoding not only preserves raw data privacy but also reduces communication overhead (Sec. 2) since the decoding based on the Mixup data augmentation can generate multiple synthetic seed samples by changing the superpositioning combinations.

Fig. 10 first verifies our conjecture that FL achieves higher accuracy when the uplink channel capacity is as high as the downlink (Uplink = Downlink). However, when the uplink channel capacity is bottlenecked (Uplink < Downlink), the accuracy of FL is significantly degraded due to its large payload sizes and the resultant frequent uploading failures within a target latency deadline. In this uplink-downlink asymmetric channel, Mix2FLD achieves higher accuracy with less variance than FL and FD.


6.9 One-Shot FL via XOR Mixup

Imbalanced data distributions could significantly degrade FL performance Sec. 3.1) [13], [36], [73]. For the MNIST and CIFAR-10 datasets wherein each worker has scarce samples of specific labels, the classification accuracy is degraded by up to 11% and 51%, respectively, compared to the IID counterparts [188]. To correct such a non-IID data problem, a straightforward solution is to exchange and fill in missing raw samples, which may however violate data privacy. Alternatively, we apply an *XOR based mixup data augmentation method (XorMixup)* that is extended to a novel one-shot FL framework, termed *XorMixFL*.

XorMixup was inspired by the Mixup data augmentation technique (Vanilla Mixup) producing a synthetic sample ($A+B$) by linearly superpositioning two raw samples A and B (Sec. 5.3) [189]. Similarly, XorMixup combines two samples not linearly but using the bit-wise XOR operation \oplus that has the following flipping property: $(A\oplus B)\oplus B = A$. To preserve the data privacy while generating realistic synthetic



(a) A schematic illustration of XorMixFL.



(b) Test accuracy of XorMixFL.

Fig. 11. XorMixFL: (a) a schematic illustration in which non-IID data distributions are corrected by XorMixup data augmentation while preserving raw data privacy; and (b) test accuracy of XorMixFL for each label, compared to Vanilla FL and standalone training.

samples, (i) each worker encodes two local samples ($A\oplus B$) that is exchanged with other devices, and (ii) the received ($A\oplus B$) is decoded not using the original B but a sample B' stored in a different worker, which has the same label of B . Consequently, the decoding yields $(A\oplus B)\oplus B' = A'$ that reflects some key features of A but is not the same as A . Owing to the mixing nature, both (i) and (ii) preserve raw data privacy across different workers, while (ii) improves the synthetic sample’s authenticity, increasing one-shot FL accuracy as elaborated next.

As illustrated in Fig. 11(a), by applying XorMixup to a one-shot FL framework having only one communication round [190], [191], each device in XorMixFL uploads its encoded seed samples to a server. The server decodes and augments the seed samples using its own base samples until all the samples are evenly distributed across labels. The server can be treated as one of the devices, or a parameter server storing an imbalanced dataset. Then, utilizing the reconstructed dataset, the server trains a global model that is downloaded by each device until convergence. Under a non-IID MNIST dataset, simulation results in Fig. 11(b)

corroborate that XorMixFL achieves up to 8.13% and 17.6% higher accuracy than standalone ML and Vanilla FL, respectively.

6.10 Tripartite SL for Medical Diagnosis


In this use case, we study a privacy-preserving SL framework (Sec. 5.1) for multiple medical platforms (e.g., hospitals or e-health wearables). These platforms store their own privacy-sensitive medical data, and are willing to cooperatively train a global model by the aid of a server storing a fraction of the model. Specifically, we consider a medical image classification task, in which not only the raw samples (e.g., a chest X-ray images) but also their ground-truth labels (e.g., lung cancer diagnosis) are privacy sensitive. In an NN model, each raw sample is fed to the input layer, and its ground-truth label is compared with the model’s prediction for loss calculation at the output layer. Therefore, to preserve the data privacy of each sample-and-label pair, both input and output layers should be stored by each platform, while the rest of the layers can be offloaded to the server, resulting in tripartite SL. This is in stark contrast to the standard bipartite SL where only the input layer is stored at each worker, while the remaining layers can be offloaded to the server.

Following the aforementioned tripartite SL, as illustrated in Fig. 12(a), we consider a single NN having k layers whose input layer L_1 and output layer L_k are stored at each platform, while the rest is run at the server. As depicted in Fig. 12(b), for each iteration, in the forward propagation, the activation of L_1 and L_{k-1} are exchanged between a platform and the server without revealing raw samples. After calculating the loss at the output layer stored at the platform, the gradients of L_k and L_2 are exchanged while hiding the ground-truth labels. While effective in preserving data privacy, the communication efficiency of tripartite SL is questionable due to frequent forward and backward propagations over wireless channels.

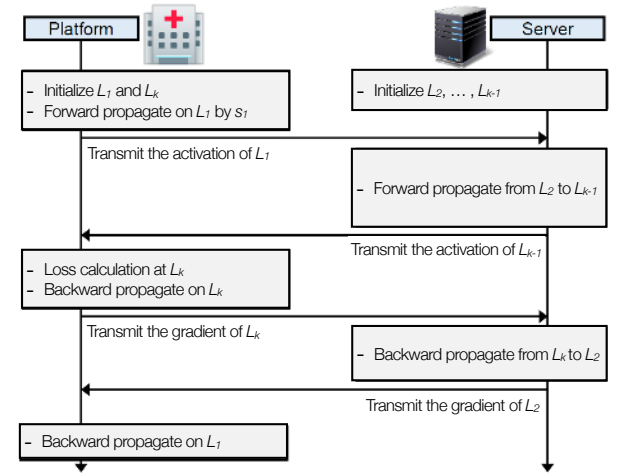
To validate its communication efficiency, we compare tripartite SL with the large-scale minibatch stochastic gradient descent (LS-SGD) [192] by measuring their transmitted data until convergence under VGG and ResNet NN model architectures with a medical X-ray dataset, CheXpert [193]. Fig. 12(c) shows that under VGG, tripartite SL yields 0.8 GB transmitted data to achieve 95% test accuracy, while LS-SGD incurs 2 GB transmitted data with 55% accuracy. A similar tendency can be observed under ResNet, in which tripartite SL consumes 0.5 GB transmitted data with 75% accuracy, whereas LS-SGD results in 1.5 GB transmitted data with 10% accuracy. This experiment concludes that in spite of more frequent communications due to the layer splits and exchanging instantaneous forward/backward propagations, rather than periodically exchanging model parameters, tripartite SL ends up with achieving lower total communication cost until the convergence. This is viable thanks to its much less communication rounds (i.e., faster convergence) and much smaller communication payload sizes.

6.11 Channel and Packet Adaptive Parallel SL

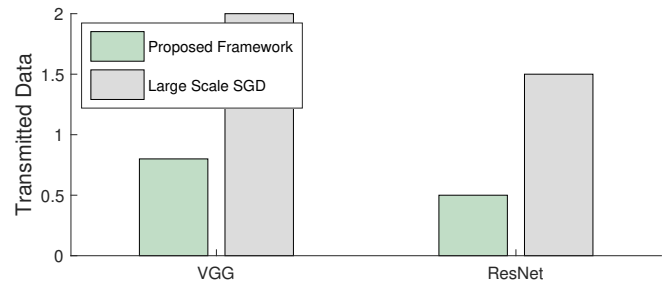
Vanilla SL is inefficient in terms of communication energy consumption when supporting multiple devices through



(a) A schematic illustration of tripartite SL.



(b) A flowchart of tripartite SL.



(c) The communication cost of tripartite SL until convergence under VGG and ResNet NN architectures, compared to large-scale minibatch SGD.

Fig. 12. Tripartite SL: (a) a schematic illustration in which the input and output layers are co-located with the data owner while the remaining hidden layers are stored at a server; (b) a flowchart clarifying forward and backward propagations; and (c) communication cost compared to the mini-batch SGD.

wireless channels (Sec. 5.1). Consider a server storing a common upper segment of NN layers that are associated with multiple devices storing its lower segments and feeding their own data samples. For these multiple devices, Vanilla SL is often implemented in a sequential manner preventing multiple devices to simultaneously connect with the server. Concatenating the output features of multiple devices into a single large vector and feeding into the server can improve the SL performance [21], with increased transmission energy consumption in the uplink and back propagation overhead in the downlink. Lastly, since the model structure cannot be dynamically adjusted during training and inference, for a fixed dimension of the server’s input layer, the server


needs to wait until the input layer is entirely filled with a predetermined number of devices or to pad arbitrary values for straggling devices due to poor channel conditions, increasing latency or degrading accuracy, respectively. Moreover, straggling devices due to intermittent connectivity under poor channel conditions (Sec. 2) either increase waiting times of acquiring the input layer at the server or persuade the server to pad arbitrary values yielding loss of accuracy.

In this view, parallel SL architecture utilizing feature averaging via Mixup augmentation-based (Section 5.3) multiple devices' outputs super-positioning can be used [149]. Adopting feature averaging in contrast to output concatenation allows server's input dimension to remain fixed independent from the number of contributing devices, enabling communication and energy efficient scalability with low training latency as illustrated in Fig. 13(a). Additionally, controlling the batch size, packet sizes of the devices' cut layer's activation to be exchanged with the server can be controlled. With small batch sizes, the accuracy can be improved with the cost of degraded uplink data rates over short packets, which can be resolved by data aggregation as discussed in Sec. 4.3. The tradeoff between test accuracy and training latency based on short packet aggregation for different choices of batch sizes is illustrated in Fig. 13(b).


6.12 Heteromodal SL for mmWave Channel Prediction

In this use case, we focus on predicting future mmWave channels by utilizing preceding mmWave signal received signal strength (RSS) history and image frames captured by two RGB depth (RGB-D) cameras mounted in different locations. Fusing these multiple modalities are essential in improving the prediction accuracy by complementing missing features one another. In particular, camera images involve useful features of blockage mobility patterns determining sudden line-of-sight (LOS) and NLOS transitions that are hardly observed from RSS, whereas RSS better describes short-term channel fluctuations for a given LOS or NLOS channel condition. Furthermore, the use of multiple cameras can overcome occlusions and missing frames (Sec. 2) due to the limited field-of-views (FoVs) and insufficient frame rates of cameras, respectively. It is however challenging to fuse such multimodal and heterogeneous data. Indeed, these data are non-IID, under which FL and its variants cannot achieve high accuracy as highlighted in Sec. 3.1.

In this regard, a joint design of SL (Sec. 5.1) that is robust against non-IID data distributions [82], [194] and feature interpolation and averaging via the Mixup data augmentation (Sec. 5.3) with heterogeneous FoVs and frame rates improving energy efficiency is considered. In the SL design, each camera feeds a sequence of image frames into its convolutional and recurrent layers whose output is uploaded to a BS's fully connected layers at which the BS's uplink mmWave RSS is fused with the uploaded features from the cameras. The proposed SL framework is validated by simulation with data measured in a real experiment using 60 GHz mmWave signals and two Kinect RGB-D cameras [195]. When predicting the future uplink mmWave RSS in 500ms by observing a sequence of RSS or image frames during 100 ms. To improve accuracy without degrading



(a) Performance comparison of Parallel SL in terms of test accuracy over training duration.



(b) Impact of batch size on the tradeoff between test accuracy and training latency in Parallel SL.

Fig. 13. Parallel SL: (a) Learning curve of parallel SL compared to sequential (Vanilla) SL, FL, and FD ; and (b) accuracy (blue, left) vs. communication latency (yellow, right) for different batch sizes.

communication efficiency, the sequence of image features generated from the camera with a lower frame rate, missing feature elements are interpolated by equally superpositioning neighboring features via manifold Mixup (Sec. 5.3). Such an interpolation reduces the non-IIDness induced by the heterogeneous frame rates, yielding higher accuracy as shown in Fig. 14(a). Note that this manifold Mixup for feature interpolation is performed within a sequence of features, whereas the aforementioned manifold Mixup for feature averaging is performed across the sequences uploaded from different cameras. Compared to a baseline scheme directly interpolating missing frames at cameras before transmissions, the aforementioned interpolation is performed at the BS after transmissions without increasing the communication payload sizes achieving low transmission latency as observed in Fig. 14(b), while yielding low power consumption at both cameras and BS as shown in Figs. 14(c) and (d).

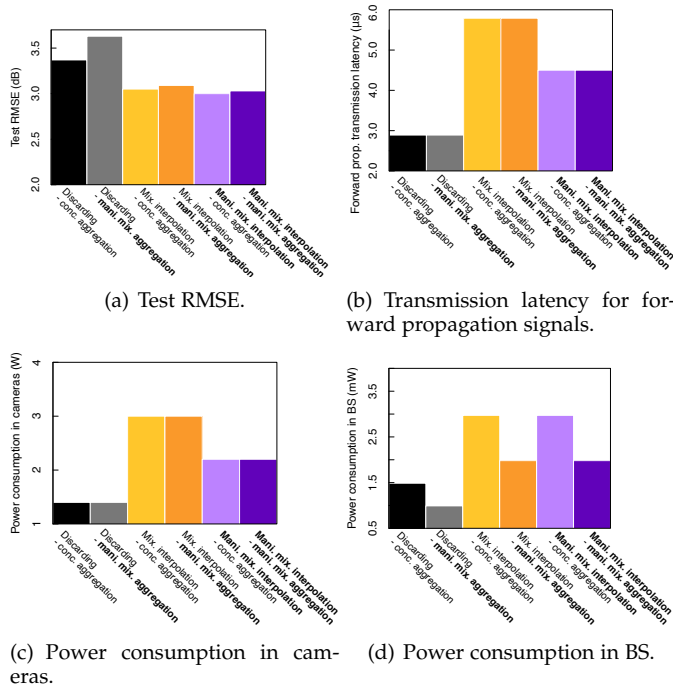


Fig. 14. Performance comparison of Heteromodal SL in terms of: (a) accuracy, (b) communication latency, (c) energy consumption at cameras (workers), and (d) energy consumption at BS (server).

7 CONCLUDING REMARKS

Imbuing intelligence into edge devices enables low-latency and scalable decision-making at the network edge in 5G communication systems and beyond. On the other hand, updating outdated edge intelligence mandates communication with federating edge devices, improving the accuracy and reliability of the decision-making at the edge. To create greater synergy, this work has explored communication-efficient and distributed learning frameworks and their use cases by co-designing ML and communication principles under various challenges incurred by communication, computing, energy, and data privacy issues. The overarching goal of this article is to foster more fundamental research in this direction and bridge connections between communication and ML communities.



Jihong Park (S'09-M'16) is a Lecturer (assistant professor) at the School of IT, Deakin University, Australia. He received the B.S. and Ph.D. degrees from Yonsei University, Seoul, Korea, in 2009 and 2016, respectively. He was a Post-Doctoral Researcher with Aalborg University, Denmark, from 2016 to 2017; the University of Oulu, Finland, from 2018 to 2019. His recent research focus includes communication-efficient distributed learning, distributed control, and distributed ledger technology, as well as their applications for beyond 5G/6G communication systems. He served as a Conference/Workshop Program Committee Member for IEEE GLOBECOM, ICC, and INFOCOM, as well as NeurIPS, ICML, and IJCAI. He received the IEEE GLOBECOM Student Travel Grant in 2014, the IEEE Seoul Section Student Paper Contest Bronze Prize in 2014, and the 6th IDIS-ETNEWS (The Electronic Times) Paper Contest Award sponsored by the Ministry of Science, ICT, and Future Planning of Korea. Currently, he is an Associate Editor of *Frontiers in Data Science for Communications*, a Review Editor of *Frontiers in Aerial and Space Networks*, a Review Editor of *Frontiers in Signal Processing for Communications*, and a Guest Editor of MDPI *Telecom SI* on "millimeter wave communications and networking in 5G and beyond."



Sumudu Samarakoon (S'08-AM'18) received his B. Sc. Degree (Hons.) in Electronic and Telecommunication Engineering from the University of Moratuwa, Sri Lanka in 2009, the M. Eng. degree from the Asian Institute of Technology, Thailand in 2011, and Ph. D. degree in Communication Engineering from University of Oulu, Finland in 2017. He is currently working in Centre for Wireless Communications, University of Oulu, Finland as a post doctoral researcher. His main research interests are in heterogeneous networks, small cells, radio resource management, reinforcement learning, and game theory. In 2016, he received the Best Paper Award at the European Wireless Conference and Excellence Awards for innovators and the outstanding doctoral student in the Radio Technology Unit, CWC, University of Oulu.



Anis Elgabri is a postdoctoral researcher at the Centre for Wireless Communications, University of Oulu. He received the B.Sc. degree in electrical and electronic engineering from the University of Tripoli, Libya, in 2004, the M.Eng. degree from UKM, Malaysia, in 2007, and MSc and PhD from the department of electrical and computer engineering, Purdue university, Indiana, USA in 2015 and 2018 respectively. His main research interests are in heterogeneous networks, radio resource management, vehicular communication, video streaming, and distributed machine learning. He was the recipient of the best paper award in HotSpot workshop, 2018 (Infocom 2018).



Mehdi Bennis is an Associate Professor at the Centre for Wireless Communications, University of Oulu, Finland, an Academy of Finland Research Fellow and head of the intelligent connectivity and networks/systems group (ICON). His main research interests are in radio resource management, heterogeneous networks, game theory and machine learning in 5G networks and beyond. He has co-authored one book and published more than 200 research papers in international conferences, journals and book chapters.

He has been the recipient of several prestigious awards including the 2015 Fred W. Ellersick Prize from the IEEE Communications Society, the 2016 Best Tutorial Prize from the IEEE Communications Society, the 2017 EURASIP Best paper Award for the Journal of Wireless Communications and Networks, the all-University of Oulu award for research and the 2019 IEEE ComSoc Radio Communications Committee Early Achievement Award. Dr Bennis is an editor of IEEE Transactions on Communications and an IEEE Fellow.



Joongheon Kim (M'06–SM'18) has been with the School of Electrical Engineering, Korea University, Seoul, Korea, since 2019, where he is currently an assistant professor. He received the B.S. and M.S. degrees in computer science and engineering from Korea University, Seoul, Korea, in 2004 and 2006, respectively; and the Ph.D. degree in computer science from the University of Southern California (USC), Los Angeles, CA, USA, in 2014. Before joining Korea University, he was with LG Electronics (Seoul, Korea, 2006–

2009), InterDigital (San Diego, CA, USA, 2012), Intel Corporation (Santa Clara in Silicon Valley, CA, USA, 2013–2016), and Chung-Ang University (Seoul, Korea, 2016–2019). He serves as an associate editor for IEEE Transactions on Vehicular Technology. He was a recipient of Annenberg Graduate Fellowship with his Ph.D. admission from USC (2009), Intel Corporation Next Generation and Standards (NGS) Division Recognition Award (2015), Haedong Young Scholar Award by the Korea Institute of Communication and Information Sciences (KICS) (2018), IEEE Vehicular Technology Society (VTS) Seoul Chapter Award (2019), Outstanding Contribution Award by KICS (2019), Gold Paper Award from IEEE Seoul Section Student Paper Contest (2019), and IEEE Systems Journal Best Paper Award (2020).



Seong-Lyun Kim is currently a Professor and Head of the School of Electrical & Electronic Engineering, Yonsei University, Seoul, Korea, leading the Robotic & Mobile Networks Laboratory (RAMO) and the Center for Flexible Radio (CFR+). He is co-directing H2020 EUK PriMO-5G project, and the chair of Smart Factory Committee of 5G Forum, Korea. He was an Assistant Professor of Radio Communication Systems at the Department of Signals, Sensors & Systems, Royal Institute of Technology (KTH), Stockholm, Sweden. He was a Visiting Professor at the Control Engineering Group, Helsinki University of Technology (now Aalto), Finland, the KTH Center for Wireless Systems, and the Graduate School of Informatics, Kyoto University, Japan. He served as a technical committee member or a chair for various conferences, and an editorial board member of IEEE Transactions on Vehicular Technology, IEEE Communications Letters, Elsevier Control Engineering Practice, Elsevier ICT Express, and Journal of Communications and Network. His research interest includes radio resource management, information theory in wireless networks, collective intelligence, and robotic networks.

Sweden. He was a Visiting Professor at the Control Engineering Group, Helsinki University of Technology (now Aalto), Finland, the KTH Center for Wireless Systems, and the Graduate School of Informatics, Kyoto University, Japan. He served as a technical committee member or a chair for various conferences, and an editorial board member of IEEE Transactions on Vehicular Technology, IEEE Communications Letters, Elsevier Control Engineering Practice, Elsevier ICT Express, and Journal of Communications and Network. His research interest includes radio resource management, information theory in wireless networks, collective intelligence, and robotic networks.



Mérouane Debbah (S'01-M'04-SM'08-F'15) received the M.Sc. and Ph.D. degrees from the Ecole Normale Supérieure Paris-Saclay, France. He was with Motorola Labs, Saclay, France, from 1999 to 2002, and also with the Vienna Research Center for Telecommunications, Vienna, Austria, until 2003. From 2003 to 2007, he was an Assistant Professor with the Mobile Communications Department, Institut Eurecom, Sophia Antipolis, France. From 2007 to 2014, he was the Director of the Alcatel-Lucent Chair on Flexible

Radio. Since 2007, he has been a Full Professor with CentraleSupélec, Gif-sur-Yvette, France. Since 2014, he has been a Vice-President of the Huawei France Research Center and the Director of the Mathematical and Algorithmic Sciences Lab. He has managed 8 EU projects and more than 24 national and international projects. His research interests lie in fundamental mathematics, algorithms, statistics, information, and communication sciences research. He is an IEEE Fellow, a WWRF Fellow, and a Membre émérite SEE. He was a recipient of the ERC Grant MORE (Advanced Mathematical Tools for Complex Network Engineering) from 2012 to 2017. He was a recipient of the Mario Boella Award in 2005, the IEEE Glavieux Prize Award in 2011, and the Qualcomm Innovation Prize Award in 2012. He received 20 best paper awards, among which the 2007 IEEE GLOBECOM Best Paper Award, the Wi-Opt 2009 Best Paper Award, the 2010 Newcom++ Best Paper Award, the WUN CogCom Best Paper 2012 and 2013 Award, the 2014 WCNC Best Paper Award, the 2015 ICC Best Paper Award, the 2015 IEEE Communications Society Leonard G. Abraham Prize, the 2015 IEEE Communications Society Fred W. Ellersick Prize, the 2016 IEEE Communications Society Best Tutorial Paper Award, the 2016 European Wireless Best Paper Award, the 2017 Eurasip Best Paper Award, the 2018 IEEE Marconi Prize Paper Award, the 2019 IEEE Communications Society Young Author Best Paper Award and the Valuetools 2007, Valuetools 2008, CrownCom 2009, Valuetools 2012, SAM 2014, and 2017 IEEE Sweden VT-COM-IT Joint Chapter best student paper awards. He is an Associate Editor-in-Chief of the journal Random Matrix: Theory and Applications. He was an Associate Area Editor and Senior Area Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING from 2011 to 2013 and from 2013 to 2014, respectively.

REFERENCES

- [1] M. Bennis, M. Debbah, and V. Poor, "Ultra-reliable and low-latency wireless communication: Tail, risk and scale," *Proceedings of the IEEE*, vol. 106, pp. 1834–1853, Oct. 2018.
- [2] J. Park, S. Samarakoon, H. Shiri, M. K. Abdel-Aziz, T. Nishio, A. Elgabri, and M. Bennis, "Extreme URLLC: Vision, challenges, and key enablers," *arXiv preprint arXiv:2001.09683*, 2020.
- [3] I. WP5D, "Minimum requirements related to technical performance for IMT-2020 radio interface(s)," 2017.
- [4] S. R. Pokhrel, J. Ding, J. Park, O.-S. Park, and J. Choi, "Towards enabling critical mMTC: A review of URLLC within mMTC," *IEEE Access*, vol. 8, pp. 131796–131813, 2020.
- [5] M. Latva-aho and K. Leppänen, "Key drivers and research challenges for 6g ubiquitous wireless intelligence," in *white paper*, University of Oulu, 2019.
- [6] Alen Space, "A basic guide to nanosatellites." [online, Accessed: 2020-07-30]. <https://alen.space/basic-guide-nanosatellites>.
- [7] STARLINK, "High speed internet access across the globe." [online, Accessed: 2020-08-02]. <https://www.starlink.com>.
- [8] Amazon, "Project Kuiper." [online, Accessed: 2020-08-02]. <https://www.amazon.jobs/en/teams/projectkuiper>.
- [9] OneWeb, "How OneWeb is changing global communications." [online, Accessed: 2020-08-02]. <https://www.oneweb.world>.
- [10] J.-H. Lee, J. Park, M. Bennis, and Y.-C. Ko, "Integrating LEO satellite and UAV relaying via reinforcement learning for non-terrestrial networks," *arXiv preprint arXiv:2005.12521*, 2020.
- [11] J.-H. Lee, J. Park, M. Bennis, and Y.-C. Ko, "Integrating leo satellites and multi-uav reinforcement learning for hybrid fso/rrf non-terrestrial networks," 2020.
- [12] M. Chen, U. Challita, W. Saad, C. Yin, and M. Debbah, "Machine learning for wireless networks with artificial intelligence: A tutorial on neural networks," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 4, pp. 3039–3071, 2019.
- [13] J. Park, S. Samarakoon, M. Bennis, and M. Debbah, "Wireless network intelligence at the edge," *Proceedings of the IEEE*, vol. 107, pp. 2204–2239, October 2019.

- [14] S. Dörner, S. Cammerer, and J. Hoydis, "Deep learning-based communication over the air," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, pp. 132–143, Feb. 2018.
- [15] X. Wang, Y. Han, C. Wang, Q. Zhao, X. Chen, and M. Chen, "In-edge AI: Intelligentizing mobile edge computing, caching and communication by federated learning," *ArXiv preprint*, vol. abs/1809.07857, Sept. 2018.
- [16] I. Ahmad, S. Shahabuddin, H. Malik, E. Harjula, T. Leppänen, L. Lovén, A. Anttonen, A. H. Sodhro, M. M. Alam, M. Juntti, A. Yla-Jaaski, T. Sauter, A. Gurtov, M. Ylianttila, and J. Riekk, "Machine learning meets communication networks: Current trends and future challenges," *IEEE Access*, pp. 1–1, 2020.
- [17] Ericsson Blog, "TinyML as a service and the challenges of machine learning at the edge." [online, Accessed: 2020-07-30]. <https://www.ericsson.com/en/blog/2019/12/tinyml-as-a-service>.
- [18] J. Konecny, H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon, "Federated learning: strategies for improving communication efficiency," in *Proc. of NIPS Wksp. PMPML*, (Barcelona, Spain), Dec. 2016.
- [19] R. Anil, G. Pereyra, A. Passos, R. Ormandi, G. E. Dahl, and G. E. Hinton, "Large scale distributed neural network training through online distillation," in *Proc. International Conference on Learning Representations (ICLR)*, Vancouver, BC, Canada, April-May 2018.
- [20] J. Ahn, O. Simeone, and J. Kang, "Wireless federated distillation for distributed edge learning with heterogeneous data," *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, Istanbul, Turkey, Sep. 2019.
- [21] P. Vepakomma, O. Gupta, T. Swedish, and R. Raskar, "Split learning for health: Distributed deep learning without sharing raw patient data," *Arxiv preprint*, vol. abs/1812.00564, Dec. 2018.
- [22] S. Haykin, *An introduction to analog and digital communication*. John Wiley, 1994.
- [23] Forbes, "The next generation of artificial intelligence." [online, Accessed: 2020-08-02], October 2020.
- [24] Atis, "Promoting U.S. leadership on the path to 6G," white paper, May 2020.
- [25] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology*, vol. 10, February 2019.
- [26] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Processing Magazine*, vol. 37, pp. 50–60, May 2020.
- [27] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al., "Advances and open problems in federated learning," *arXiv preprint arXiv:1912.04977*, 2019.
- [28] Y. Shi, K. Yang, T. Jiang, J. Zhang, and K. B. Letaief, "Communication-efficient edge ai: Algorithms and systems," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 4, pp. 2167–2191, 2020.
- [29] M. Chen, H. V. Poor, W. Saad, and S. Cui, "Wireless communications for collaborative federated learning," 2020.
- [30] C. Briggs, Z. Fan, and P. Andras, "A review of privacy-preserving federated learning for the internet-of-things," 2020.
- [31] L. Lyu, H. Yu, and Q. Yang, "Threats to federated learning: A survey," 2020.
- [32] Y. Liu, J. Peng, J. Kang, A. M. Iliyasu, D. Niyato, and A. A. A. El-Latif, "A secure federated learning framework for 5g networks," *IEEE Wireless Communications*, vol. 27, no. 4, pp. 24–31, 2020.
- [33] P. Vepakomma, T. Swedish, R. Raskar, O. Gupta, and A. Dubey, "No peek: A survey of private distributed deep learning," 2018.
- [34] C. He, S. Li, J. So, M. Zhang, H. Wang, X. Wang, P. Vepakomma, A. Singh, H. Qiu, L. Shen, P. Zhao, Y. Kang, Y. Liu, R. Raskar, Q. Yang, M. Annavaram, and S. Avestimehr, "Fedml: A research library and benchmark for federated machine learning," *arXiv preprint arXiv:2007.13518*, 2020.
- [35] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of machine learning*. MIT press, 2018.
- [36] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated Learning with Non-IID Data," [Online]: arXiv preprint arXiv:1806.00582, 2018.
- [37] M. Frid-Adar, I. Diamant, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan, "GAN-based synthetic medical image augmentation for increased CNN performance in liver lesion classification," *Neurocomputing*, vol. 321, pp. 321–331, 2018.
- [38] C. Bowles, L. Chen, R. Guerrero, P. Bentley, R. Gunn, A. Hammers, D. A. Dickie, M. V. Hernández, J. Wardlaw, and D. Rueckert, "Gan augmentation: Augmenting training data using generative adversarial networks," *arXiv preprint arXiv:1810.10863*, 2018.
- [39] W.-N. Hsu, Y. Zhang, and J. Glass, "Unsupervised domain adaptation for robust speech recognition via variational autoencoder-based data augmentation," in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 16–23, IEEE, 2017.
- [40] H. Nishizaki, "Data augmentation and feature extraction using variational autoencoder for acoustic modeling," in *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pp. 1222–1227, IEEE, 2017.
- [41] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pp. 1322–1333, 2015.
- [42] S. Ahn, J. Kim, E. Lim, W. Choi, A. Mohaisen, and S. Kang, "Shm-Caffe: a distributed deep learning platform with shared memory buffer for HPC architecture," in *Proc. 2018 IEEE International Conference on Distributed Computing Systems (ICDCS)*, pp. 1118–1128, 2018.
- [43] K. Lee, "Nvidia GeForce RTX 2080 Ti review." [online, Accessed: 2020-07-30]. <https://www.techradar.com/reviews/nvidia-geforce-rtx-2080-ti-review>.
- [44] N. Wang, C.-Y. Chen, and K. Gopalakrishnan, "Ultra-low-precision training of deep neural networks." [online, Accessed: 2020-08-02]. <https://www.ibm.com/blogs/research/2019/05/ultra-low-precision-training/>.
- [45] A. Rodriguez, "Lowering numerical precision to increase deep learning performance." [online, Accessed: 2020-08-02]. <https://www.intel.com/content/www/us/en/artificial-intelligence/posts/lowering-numerical-precision-increase-deep-learning-performance.html>.
- [46] T. Murovič and A. Trost, "Massively parallel combinational binary neural networks for edge processing," *Elektrotehnicki Vestnik*, vol. 86, no. 1/2, pp. 47–53, 2019.
- [47] J. Yang, X. Ge, and Y. Zhong, "How much of wireless rates can smartphones support in 5g networks?," *IEEE Network*, vol. 33, no. 3, pp. 122–129, 2019.
- [48] X. Wang, L. Kong, F. Kong, F. Qiu, M. Xia, S. Arnon, and G. Chen, "Millimeter wave communication: A comprehensive survey," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, pp. 1616–1653, 2018.
- [49] D. S. Lun, M. Médard, R. Koetter, and M. Effros, "On coding for reliable communication over packet networks," *Physical Communication*, vol. 1, no. 1, pp. 3–20, 2008.
- [50] T. Yang, G. Andrew, H. Eichner, H. Sun, W. Li, N. Kong, D. Ramage, and F. Beaufays, "Applied federated learning: Improving google keyboard query suggestions," *arXiv preprint arXiv:1812.02903*, 2018.
- [51] S. Samarakoon, M. Bennis, W. Saad, and M. Debbah, "Distributed federated learning for ultra-reliable low-latency vehicular communications," *IEEE Transactions on Communications*, vol. 68, no. 2, pp. 1146–1159, 2019.
- [52] J. Konečný, B. McMahan, and D. Ramage, "Federated optimization: Distributed optimization beyond the datacenter," *arXiv preprint arXiv:1511.03575*, 2015.
- [53] V. Smith, C.-K. Chiang, M. Sanjabi, and A. S. Talwalkar, "Federated multi-task learning," in *Proc. of NIPS (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds.)*, (Long Beach, USA), pp. 4424–4434, Curran Associates, Inc., Dec. 2017.
- [54] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-IID data," *ArXiv preprint*, vol. abs/1806.00582, June 2018.
- [55] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, "Robust and communication-efficient federated learning from non-iid data," *IEEE transactions on neural networks and learning systems*, 2019.
- [56] A. Lalitha, O. C. Kilinc, T. Javidi, and F. Koushanfar, "Peer-to-peer federated learning on graphs," *arXiv preprint arXiv:1901.11173*, 2019.
- [57] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, (New York, NY, USA), pp. 1322–1333, Association for Computing Machinery, 2015.

- [58] B. Hitaj, G. Ateniese, and F. Perez-Cruz, "Deep models under the gan: Information leakage from collaborative deep learning," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*, (New York, NY, USA), pp. 603–618, Association for Computing Machinery, 2017.
- [59] Y. Zhao, J. Zhao, L. Jiang, R. Tan, D. Niyato, Z. Li, L. Lyu, and Y. Liu, "Privacy-preserving blockchain-based federated learning for iot devices," *IEEE Internet of Things Journal*, pp. 1–1, 2020.
- [60] Y. Zhao, J. Zhao, M. Yang, T. Wang, N. Wang, L. Lyu, D. Niyato, and K. Y. Lam, "Local differential privacy based federated learning for internet of things," 2020.
- [61] M. Yang, L. Lyu, J. Zhao, T. Zhu, and K.-Y. Lam, "Local differential privacy and its applications: A comprehensive survey," 2020.
- [62] Y. Koda, K. Yamamoto, T. Nishio, and M. Morikura, "Differentially private aircomp federated learning with power adaptation harnessing receiver noise," 2020.
- [63] A. Elgabli, J. Park, C. B. Issaid, and M. Bennis, "Harnessing wireless channels for scalable and privacy-preserving federated learning," *arXiv preprint arXiv:2007.01790*, 2020.
- [64] H. Kim, J. Park, M. Bennis, and S.-L. Kim, "Blockchained on-device federated learning," to appear in *IEEE Communications Letters* [Online]. *ArXiv preprint: abs/1808.03949*.
- [65] A. Elgabli, J. Park, A. S. Bedi, M. Bennis, and V. Aggarwal, "GADMM: Fast and communication efficient framework for distributed machine learning," *Journal of Machine Learning Research*, vol. 21, no. 76, pp. 1–39, 2020.
- [66] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., "Language models are few-shot learners," *arXiv preprint arXiv:2005.14165*, 2020.
- [67] G. E. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," in *Proc. of NIPS Wksp. Deep Learning*, (Montréal, Canada), pp. 1–9, Dec. 2014.
- [68] J. Wang, W. Bao, L. Sun, X. Zhu, B. Cao, and P. S. Yu, "Private model compression via knowledge distillation," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 1190–1197, Jul. 2019.
- [69] P. Han, J. Park, S. Wang, and Y. Liu, "Robustness and diversity seeking data-free knowledge distillation," 2020.
- [70] L. Lyu and C.-H. Chen, "Differentially private knowledge distillation for mobile analytics," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20*, (New York, NY, USA), pp. 1809–1812, Association for Computing Machinery, 2020.
- [71] H. Seo, J. Park, S. Oh, M. Bennis, and S.-L. Kim, "Federated knowledge distillation," 2020.
- [72] H. Cha, J. Park, H. Kim, M. Bennis, and S.-L. Kim, "Federated reinforcement distillation with proxy experience replay memory," to appear in *IEEE Intelligent Systems*.
- [73] S. Oh, J. Park, E. Jeong, H. Kim, M. Bennis, and S.-L. Kim, "Mix2FLD: Downlink federated learning after uplink federated distillation with two-way mixup," to appear in *IEEE Communications Letters*.
- [74] J. Ahn, O. Simeone, and J. Kang, "Cooperative learning via federated distillation over fading channels," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Barcelona, Spain, May 2020*.
- [75] P. Vepakomma, O. G. A. Dubey, and R. Raskar, "Reducing leakage in distributed deep learning for sensitive health data," in *Proc. of ICLR*, (New Orleans, USA), May 2019.
- [76] L. Lyu, J. C. Bezdek, J. Jin, and Y. Yang, "Foreseen: Towards differentially private deep inference for intelligent internet of things," *IEEE Journal on Selected Areas in Communications*, vol. 38, no. 10, pp. 2418–2429, 2020.
- [77] L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov, "Exploiting unintended feature leakage in collaborative learning," in *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 691–706, 2019.
- [78] L. Lyu, Y. Li, X. He, and T. Xiao, "Towards differentially private text representations," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '20*, (New York, NY, USA), pp. 1813–1816, Association for Computing Machinery, 2020.
- [79] J. Jeon, J. Kim, J. Kim, K. Kim, A. Mohaisen, and J.-K. Kim, "Privacy-preserving deep learning computation for geodistributed medical big-data platforms," in *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks—Supplemental Volume (DSN-S)*, pp. 3–4, IEEE, 2019.
- [80] Y. Koda, J. Park, M. Bennis, K. Yamamoto, T. Nishio, and M. Morikura, "One pixel image and RF signal based split learning for mmWave received power prediction," in *Proceedings of the 15th International Conference on emerging Networking EXperiments and Technologies*, pp. 54–56, 2019.
- [81] Y. Koda, J. Park, M. Bennis, K. Yamamoto, T. Nishio, M. Morikura, and K. Nakashima, "Communication-efficient multimodal split learning for mmWave received power prediction," *IEEE Communications Letters*, vol. 24, no. 6, pp. 1284–1288, 2020.
- [82] Y. Koda, J. Park, M. Bennis, K. Yamamoto, T. Nishio, and M. Morikura, "Distributed heteromodal split learning for vision aided mmWave received power prediction," *arXiv preprint arXiv:2007.08208*, 2020.
- [83] A. Singh, P. Vepakomma, O. Gupta, and R. Raskar, "Detailed comparison of communication efficiency of split learning and federated learning," *arXiv preprint arXiv:1909.09145*, 2019.
- [84] K. Zhang, Z. Yang, and T. Başar, "Multi-agent reinforcement learning: A selective overview of theories and algorithms," *arXiv preprint arXiv:1911.10635*, 2019.
- [85] A. Khan, C. Zhang, D. D. Lee, V. Kumar, and A. Ribeiro, "Scalable centralized deep multi-agent reinforcement learning via policy gradients," *arXiv preprint arXiv:1805.08776*, 2018.
- [86] X. Wang and T. Sandholm, "Reinforcement learning to play an optimal Nash equilibrium in team Markov games," in *Advances in neural information processing systems*, pp. 1603–1610, 2003.
- [87] T. Chen, G. Giannakis, T. Sun, and W. Yin, "LAG: Lazily aggregated gradient for communication-efficient distributed learning," in *Advances in Neural Information Processing Systems*, pp. 5050–5060, 2018.
- [88] J. C. Duchi, A. Agarwal, and M. J. Wainwright, "Dual averaging for distributed optimization: Convergence analysis and network scaling," *IEEE Transactions on Automatic control*, vol. 57, no. 3, pp. 592–606, 2011.
- [89] P. Jiang and G. Agrawal, "A linear speedup analysis of distributed deep learning with sparse and quantized communication," in *Advances in Neural Information Processing Systems*, pp. 2525–2536, 2018.
- [90] P. Han, S. Wang, and K. K. Leung, "Adaptive gradient sparsification for efficient federated learning: An online learning approach," *arXiv preprint arXiv:2001.04756*, 2020.
- [91] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "Qsgd: Communication-efficient sgd via gradient quantization and encoding," in *Advances in Neural Information Processing Systems*, pp. 1709–1720, 2017.
- [92] S. Magnússon, H. Shokri-Ghadikolaei, and N. Li, "On maintaining linear convergence of distributed learning and optimization under limited communication," in *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, pp. 432–436, IEEE, 2019.
- [93] J. Bernstein, Y.-X. Wang, K. Azizzadenesheli, and A. Anandkumar, "signsgd: Compressed optimisation for non-convex problems," *arXiv preprint arXiv:1802.04434*, 2018.
- [94] K. Mishchenko, E. Gorbunov, M. Takáč, and P. Richtárik, "Distributed learning with compressed gradient differences," *arXiv preprint arXiv:1901.09269*, 2019.
- [95] J. Wu, W. Huang, J. Huang, and T. Zhang, "Error compensated quantized sgd and its applications to large-scale distributed optimization," *arXiv preprint arXiv:1806.08054*, 2018.
- [96] H. Zhang, J. Li, K. Kara, D. Alistarh, J. Liu, and C. Zhang, "Zipml: Training linear models with end-to-end low precision, and a little bit of deep learning," in *International Conference on Machine Learning*, pp. 4035–4043, 2017.
- [97] W. Wen, C. Xu, F. Yan, C. Wu, Y. Wang, Y. Chen, and H. Li, "Terngrad: Ternary gradients to reduce communication in distributed deep learning," in *Advances in neural information processing systems*, pp. 1509–1519, 2017.
- [98] A. Elgabli, J. Park, A. S. Bedi, C. B. Issaid, M. Bennis, and V. Aggarwal, "Q-gadmm: Quantized group ADMM for communication efficient decentralized machine learning," 2019.
- [99] R. Zamir and M. Feder, "On universal quantization by randomized uniform/lattice quantizers," *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 428–436, 1992.
- [100] N. Shlezinger, M. Chen, Y. C. Eldar, H. V. Poor, and S. Cui, "Uveqfed: Universal vector quantization for federated learning," *arXiv preprint arXiv:2006.03262*, 2020.
- [101] F. Haddadpour, M. M. Kamani, A. Mokhtari, and M. Mahdavi, "Federated learning with compression: Unified analysis and sharp guarantees," *arXiv preprint arXiv:2007.01154*, 2020.

- [102] G. Durisi, T. Koch, and P. Popovski, "Toward massive, ultra-reliable, and low-latency wireless communication with short packets," *Proceedings of the IEEE*, vol. 104, no. 9, pp. 1711–1726, 2016.
- [103] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Channel coding rate in the finite blocklength regime," *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2307–2359, 2010.
- [104] P. Popovski, J. J. Nielsen, C. Stefanovic, E. De Carvalho, E. Strom, K. F. Trillingsgaard, A.-S. Bana, D. M. Kim, R. Kotaba, J. Park, et al., "Wireless access for ultra-reliable low-latency communication: Principles and building blocks," *Ieee Network*, vol. 32, no. 2, pp. 16–23, 2018.
- [105] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. T. P. Tang, "On large-batch training for deep learning: Generalization gap and sharp minima," *arXiv preprint arXiv:1609.04836*, 2016.
- [106] T. Nishio and R. Yonetani, "Client selection for federated learning with heterogeneous resources in mobile edge," *In Proc. Int'l Conf. Commun. (ICC), Shanghai, China*, May 2019.
- [107] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," *IEEE Journal on Selected Areas in Communications*, vol. 37, pp. 1205–1221, Jun. 2019.
- [108] H. H. Yang, Z. Liu, T. Q. S. Quek, and H. V. Poor, "Scheduling policies for federated learning in wireless networks," *arXiv preprint arXiv:1908.06287*, 2019.
- [109] M. Chen, Z. Yang, W. Saad, C. Yin, H. V. Poor, and S. Cui, "A joint learning and communications framework for federated learning over wireless networks," *arXiv preprint arXiv:1909.07972*, 20019.
- [110] M. M. Amiri and D. Gunduz, "Over-the-air machine learning at the wireless edge," *Proc. IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPWAC), Cannes, France*, July 2019.
- [111] G. Zhu, Y. Wang, and K. Huang, "Broadband analog aggregation for low-latency federated edge learning," *arXiv preprint arXiv:1812.11494*.
- [112] T. Sery and K. Cohen, "On analog gradient descent learning over multiple access fading channels," *arXiv preprint arXiv:1908.07463*.
- [113] G. Zhu, Y. Du, D. Dunduz, and K. Huang, "One-bit over-the-air aggregation for communication-efficient federated edge learning: Design and convergence analysis," *arXiv preprint arXiv:2001.05713*.
- [114] M. M. Amiri and D. Gündüz, "Federated learning over wireless fading channels," *IEEE Transactions on Wireless Communications*, vol. 19, no. 5, pp. 3546–3557, 2020.
- [115] K. Yang, T. Jiang, Y. Shi, and Z. Ding, "Federated learning via over-the-air computation," *IEEE Transactions on Wireless Communications*, vol. 19, no. 3, pp. 2022–2035, 2020.
- [116] D. Liu and O. Simeone, "Privacy for free: Wireless federated learning via uncoded transmission with adaptive power control," 2020.
- [117] G. Zhu, J. Xu, K. Huang, and S. Cui, "Over-the-air computing for wireless data aggregation in massive iot," 2020.
- [118] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [119] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. of AISTATS*, (Fort Lauderdale, FL, USA), Apr. 2017.
- [120] M. M. Amiri and D. Gündüz, "Computation scheduling for distributed machine learning with straggling workers," *IEEE Transactions on Signal Processing*, vol. 67, no. 24, pp. 6270–6284, 2019.
- [121] M. S. Elbambay, C. Perfecto, C. Liu, J. Park, S. Samarakoon, X. Chen, and M. Bennis, "Wireless edge computing with latency and reliability guarantees," *Proceedings of the IEEE*, vol. 107, pp. 1717–1737, Aug. 2019.
- [122] M. Polese, R. Jana, V. Kounev, K. Zhang, S. Deb, and M. Zorzi, "Machine learning at the edge: A data-driven architecture with applications to 5G cellular networks," *IEEE Transactions on Mobile Computing*, 2020.
- [123] Y. Zhang, B. Di, P. Wang, J. Lin, and L. Song, "HetMEC: Heterogeneous multi-layer mobile edge computing in the 6 G era," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 4, pp. 4388–4400, 2020.
- [124] S. Li, M. A. Maddah-Ali, and A. S. Avestimehr, "Communication-aware computing for edge processing," in *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 2885–2889, IEEE, 2017.
- [125] T. Chanyour, M. El Ghmary, Y. Hmimz, and M. O. Cherkaoui Malki, "Energy-efficient and delay-aware multitask offloading for mobile edge computing networks," *Transactions on Emerging Telecommunications Technologies*, p. e3673, 2019.
- [126] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 6, pp. 1205–1221, 2019.
- [127] U. Mohammad and S. Sorour, "Adaptive task allocation for mobile edge learning," in *2019 IEEE Wireless Communications and Networking Conference Workshop (WCNCW)*, pp. 1–6, IEEE, 2019.
- [128] R. Xu, B. Palanisamy, and J. Joshi, "QueryGuard: Privacy-preserving latency-aware query optimization for edge computing," in *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/12th IEEE International Conference On Big Data Science And Engineering (Trust-Com/BigDataSE)*, pp. 1097–1106, IEEE, 2018.
- [129] H. H. Yang, Z. Liu, T. Q. Quek, and H. V. Poor, "Scheduling policies for federated learning in wireless networks," *IEEE Transactions on Communications*, vol. 68, no. 1, pp. 317–333, 2019.
- [130] M. M. Wadu, S. Samarakoon, and M. Bennis, "Federated learning under channel uncertainty: Joint client scheduling and resource allocation," in *2020 IEEE Wireless Communications and Networking Conference (WCNC)*, pp. 1–6, 2020.
- [131] K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konecny, S. Mazzocchi, H. B. McMahan, et al., "Towards federated learning at scale: System design," *arXiv preprint arXiv:1902.01046*, 2019.
- [132] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of fedavg on non-iid data," *arXiv preprint arXiv:1907.02189*, 2019.
- [133] H. Wang, Z. Kaplan, D. Niu, and B. Li, "Optimizing federated learning on non-iid data with reinforcement learning," in *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*, pp. 1698–1707, IEEE, 2020.
- [134] Y. Jin, L. Jiao, Z. Qian, S. Zhang, S. Lu, and X. Wang, "Resource-efficient and convergence-preserving online participant selection in federated learning,"
- [135] M. S. H. Abad, E. Ozfatura, D. Gunduz, and O. Ercetin, "Hierarchical federated learning across heterogeneous cellular networks," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8866–8870, IEEE, 2020.
- [136] Y. Jiang, S. Wang, B. J. Ko, W.-H. Lee, and L. Tassiulas, "Model pruning enables efficient federated learning on edge devices," *arXiv preprint arXiv:1909.12326*, 2019.
- [137] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," in *Advances in neural information processing systems*, pp. 1135–1143, 2015.
- [138] O. Yadan, K. Adams, Y. Taigman, and M. Ranzato, "Multi-gpu training of convnets," *arXiv preprint arXiv:1312.5853*, 2013.
- [139] Y. Kang, J. Hauswald, C. Gao, A. Rovinski, T. Mudge, J. Mars, and L. Tang, "Neurosurgeon: Collaborative intelligence between the cloud and mobile edge," *ACM SIGARCH Computer Architecture News*, vol. 45, no. 1, pp. 615–629, 2017.
- [140] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- [141] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- [142] D. P. Kingma and M. Welling, "An introduction to variational autoencoders," *Foundations and Trends in Machine Learning*, vol. 12, no. 4, pp. 307–392, 2019.
- [143] E. Jeong, S. Oh, H. Kim, J. Park, M. Bennis, and S.-L. Kim, "Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data," *arXiv preprint arXiv:1811.11479*, 2018.
- [144] E. Jeong, S. Oh, J. Park, H. Kim, M. Bennis, and S.-L. Kim, "Multi-hop federated private data augmentation with sample compression," *arXiv preprint arXiv:1907.06426*, 2019.
- [145] J. Jeon, J. Kim, J. Kim, K. Kim, A. Mohaisen, and J. Kim, "Privacy-preserving deep learning computation for geo-distributed medical big-data platforms," in *Proc. IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pp. 3–4, 2019.

- [146] M. Phuong and C. H. Lampert, "Towards understanding knowledge distillation," in *International Conference on Machine Learning (ICML), Long Beach, California*, (Long Beach, CA, USA), June 2019.
- [147] J.-H. Ahn, O. Simeone, and J. Kang, "Wireless federated distillation for distributed edge learning with heterogeneous data," in *2019 IEEE 30th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pp. 1–6, IEEE, 2019.
- [148] H. Cha, J. Park, H. Kim, S.-L. Kim, and M. Bennis, "Federated reinforcement distillation with proxy experience memory," *arXiv preprint arXiv:1907.06536*, 2019.
- [149] V. Verma, A. Lamb, C. Beckham, A. Najafi, I. Mitliagkas, D. Lopez-Paz, and Y. Bengio, "Manifold mixup: Better representations by interpolating hidden states," in *International Conference on Machine Learning*, pp. 6438–6447, 2019.
- [150] J. Park, S. Wang, A. Elgabli, S. Oh, E. Jeong, H. Cha, H. Kim, S.-L. Kim, and M. Bennis, "Distilling on-device intelligence at the network edge," [Online]. Arxiv preprint: <http://arxiv.org/abs/1907.02745>, Dec. 2018.
- [151] E. Jeong, S. Oh, H. Kim, J. Park, M. Bennis, and S.-L. Kim, "XOR Mixup: Privacy-preserving data augmentation for one-shot federated learning," *presented at 2020 ICML Wksp. Federated Learning for User Privacy and Data Confidentiality (ICML-FL)*, May 2020.
- [152] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer School on Machine Learning*, pp. 63–71, Springer, 2003.
- [153] A. Girard, C. E. Rasmussen, J. Q. Candela, and R. Murray-Smith, "Gaussian process priors with uncertain inputs application to multiple-step ahead time series forecasting," in *Advances in neural information processing systems*, pp. 545–552, 2003.
- [154] D. M. Kreps, "Nash equilibrium," in *Game Theory*, pp. 167–177, Springer, 1989.
- [155] J. Park, S. Jung, S.-L. Kim, M. Bennis, and M. Debbah, "User-centric mobility management in ultra-dense cellular networks under spatio-temporal dynamics," in *Proc. of IEEE GLOBECOM*, (Washington, DC, USA), Dec. 2016.
- [156] H. Kim, J. Park, M. Bennis, S. Kim, and M. Debbah, "Ultra-dense edge caching under spatio-temporal demand and network dynamics," in *Proc. of IEEE ICC*, (Paris, France), pp. 1–7, May 2017.
- [157] H. Kim, J. Park, M. Bennis, and S.-L. Kim, "Massive UAV-to-ground communication and its stable movement control: A mean-field approach," in *Proc. of IEEE SPAWC*, (Kalamata, Greece), June 2018.
- [158] H. Shiri, J. Park, and M. Bennis, "Massive autonomous uav path planning: A neural network based mean-field game theoretic approach," in *2019 IEEE Global Communications Conference (GLOBECOM)*, pp. 1–6, IEEE, 2019.
- [159] H. Shiri, J. Park, and M. Bennis, "Communication-efficient massive uav online path control: Federated learning meets mean-field game theory," *arXiv preprint arXiv:2003.04451*, 2020.
- [160] J.-M. Lasry and P.-L. Lions, "Mean field games," *Japanese journal of mathematics*, vol. 2, no. 1, pp. 229–260, 2007.
- [161] R. Courant, K. Friedrichs, and H. Lewy, "On the partial difference equations of mathematical physics," *IBM journal of Research and Development*, vol. 11, no. 2, pp. 215–234, 1967.
- [162] E. Gorbunov, F. Hanzely, and P. Richtárik, "Local sgd: Unified theory and new efficient methods," 2020.
- [163] A. Khaled, K. Mishchenko, and P. Richtárik, "Tighter theory for local sgd on identical and heterogeneous data," in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics* (S. Chiappa and R. Calandra, eds.), vol. 108 of *Proceedings of Machine Learning Research*, (Online), pp. 4519–4529, PMLR, 26–28 Aug 2020.
- [164] B. Woodworth, K. K. Patel, S. U. Stich, Z. Dai, B. Bullins, H. B. McMahan, O. Shamir, and N. Srebro, "Is local sgd better than minibatch sgd?," 2020.
- [165] B. Woodworth, K. K. Patel, and N. Srebro, "Minibatch vs local sgd for heterogeneous distributed learning," 2020.
- [166] A. Cotter, O. Shamir, N. Srebro, and K. Sridharan, "Better minibatch algorithms via accelerated gradient methods," 2011.
- [167] O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao, "Optimal distributed online prediction using mini-batches," 2012.
- [168] H. Yu, S. Yang, and S. Zhu, "Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning," 2018.
- [169] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, et al., "Communication-efficient learning of deep networks from decentralized data," *arXiv preprint arXiv:1602.05629*, 2016.
- [170] S. Horváth, D. Kovalev, K. Mishchenko, S. Stich, and P. Richtárik, "Stochastic distributed learning with gradient quantization and variance reduction," *arXiv preprint arXiv:1904.05115*, 2019.
- [171] D. Kovalev, A. Koloskova, M. Jaggi, P. Richtárik, and S. U. Stich, "A linearly convergent algorithm for decentralized optimization: Sending less bits for free!," 2020.
- [172] C. B. Issaid, A. Elgabli, J. Park, and M. Bennis, "Communication efficient distributed learning with censored, quantized, and generalized group admm," 2020.
- [173] M. Chen, H. V. Poor, W. Saad, and S. Cui, "Convergence time minimization of federated learning over wireless networks," in *ICC 2020 - 2020 IEEE International Conference on Communications (ICC)*, pp. 1–6, 2020.
- [174] Z. Yang, M. Chen, W. Saad, C. S. Hong, and M. Shikh-Bahaei, "Energy efficient federated learning over wireless communication networks," 2020.
- [175] A. Rahbar, A. Panahi, C. Bhattacharyya, D. Dubhashi, and M. H. Chehreghani, "On the unreasonable effectiveness of knowledge distillation: Analysis in the kernel regime," *ArXiv preprint, arXiv:2003.13438*, 2020.
- [176] A. Jacot, F. Gabriel, and C. Hongler, "Neural tangent kernel: Convergence and generalization in neural networks," in *Proc. Advances in neural information processing systems (NeurIPS)*, Montreal, Canada, December 2018.
- [177] Z. Li, D. Kovalev, X. Qian, and P. Richtárik, "Acceleration for compressed gradient descent in distributed and federated optimization," *arXiv preprint arXiv:2002.11364*, 2020.
- [178] A. Nedić, A. Olshevsky, and M. G. Rabbat, "Network topology and communication-computation tradeoffs in decentralized optimization," *Proceedings of the IEEE*, vol. 106, no. 5, pp. 953–976, 2018.
- [179] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028*, 2014.
- [180] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, "Quantum approximate optimization algorithm: performance, mechanism, and implementation on near-term devices," *arXiv preprint arXiv:1812.01041*, 2018.
- [181] J. Kim, G. Caire, and A. F. Molisch, "Quality-aware streaming and scheduling for device-to-device video delivery," *IEEE/ACM Transactions on Networking*, vol. 24, no. 4, pp. 2319–2331, 2016.
- [182] X. Ma, H. Sun, and R. Q. Hu, "Scheduling policy and power allocation for federated learning in NOMA based MEC," *ArXiv*, vol. abs/2006.13044, 2020.
- [183] J. Choi, S. Oh, and J. Kim, "Quantum approximation for wireless scheduling," *Applied Sciences*, vol. 10, no. 20, 2020.
- [184] M. Broughton, G. Verdon, T. McCourt, A. J. Martinez, J. H. Yoo, et al., "TensorFlow quantum: A software framework for quantum machine learning," *arXiv:2003.02989*, 2020.
- [185] M. Karaca, T. Alpcan, and O. Ercetin, "Smart scheduling and feedback allocation over non-stationary wireless channels," in *2012 IEEE International Conference on Communications (ICC)*, pp. 6586–6590, IEEE, 2012.
- [186] A. Bensoussan and J. F. nad Phillip Yam, *Mean Field Games and Mean Field Type Control Theory*. SpringerBriefs in Mathematics, Springer-Verlag New York, 1 ed., 2013.
- [187] J. Park, S.-L. Kim, and J. Zander, "Tractable resource management with uplink decoupled millimeter-wave overlay in ultra-dense cellular networks," *IEEE Transactions on Wireless Communications*, vol. 15, pp. 4362–4379, June 2016.
- [188] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-iid data," *arXiv preprint arXiv:1806.00582*, 2018.
- [189] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond Empirical Risk Minimization," in *Proc. of 6th International Conference on Learning Representations (ICLR)*, 2018.
- [190] N. Guha, A. Talwalkar, and V. Smith, "One-shot federated learning," [Online]. Arxiv preprint: <http://arxiv.org/abs/1902.11175>.
- [191] N. Yoshida, T. Nishio, M. Morikura, K. Yamamoto, and R. Yone-tani, "Hybrid-fl: Cooperative learning mechanism using non-iid data in wireless networks," *arXiv preprint arXiv:1905.07210*, 2019.
- [192] J. Chen, R. Monga, S. Bengio, and R. Józefowicz, "Revisiting distributed synchronous SGD," *CoRR*, vol. abs/1604.00981, 2016.
- [193] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghgo, R. L. Ball, K. S. Shpanskaya, J. Seekins,

- D. A. Mong, S. S. Halabi, J. K. Sandberg, R. Jones, D. B. Larson, C. P. Langlotz, B. N. Patel, M. P. Lungren, and A. Y. Ng, "Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison," *CoRR*, vol. abs/1901.07031, 2019.
- [194] Y. Koda, K. Yamamoto, T. Nishio, and M. Morikura, "Differentially private aircomp federated learning with power adaptation harnessing receiver noise," [Online]. Arxiv preprint: <https://arxiv.org/abs/2004.06337>.
- [195] K. Khoshelham and S. Oude Elberink, "Accuracy and resolution of kinect depth data for indoor mapping applications," *Sensors (US)*, vol. 12, no. 2, pp. 1437–1454, 2012.