

ASYNCHRONOUS DIFFUSION LEARNING WITH AGENT SUBSAMPLING AND LOCAL UPDATES

Elsa Rizk, Kun Yuan, Ali H. Sayed

ABSTRACT

In this work, we examine a network of agents operating asynchronously, aiming to discover an ideal global model that suits individual local datasets. Our assumption is that each agent independently chooses when to participate throughout the algorithm and the specific subset of its neighbourhood with which it will cooperate at any given moment. When an agent chooses to take part, it undergoes multiple local updates before conveying its outcomes to the sub-sampled neighbourhood. Under this setup, we prove that the resulting asynchronous diffusion strategy is stable in the mean-square error sense and provide performance guarantees specifically for the federated learning setting. We illustrate the findings with numerical simulations.

Index Terms— distributed systems, diffusion learning, asynchronous network, federated learning

1. INTRODUCTION AND RELATED MATERIAL

Networks of agents comprise individual agents collaborating in the pursuit of a common goal. In the case of distributed optimization, the agents aim to solve a global optimization problem by utilizing local information. Typically, the available solutions involve a local update step followed by an aggregation step among neighbours. Examples of such strategies include incremental methods [1–5], consensus methods [6–9], and diffusion methods [10–14]. Commonly, these methods assume full agent participation in every iteration and a one-to-one ratio of aggregation steps to local update steps. Nonetheless, various situations like agent drop-outs do not satisfy these assumptions [15, 16], where certain agents might not partake in each iteration. Furthermore, computational and communication limitations [17, 18] could necessitate an agent to perform several local updates before sharing outcomes with a subset of its neighbours, as opposed to all of them.

Therefore, this work focuses on asynchronous networks of agents and modifies the algorithms, particularly diffusion-type algorithms, to allow flexibility in agent participation, agent subsampling, and local updates. Relevant research on asynchronous distributed learning can be found in [19–29]. The primary relevant prior work [26] in the context of asynchronous networks operates under the assumption of independent step-size and combination weights as well as a common optimal model. However, within the scenarios we are investigating, such as federated learning [30], such assumptions do not hold. As a result, we abandon the independence assumption and study the stability and performance of a particular asynchronous distributed setup. Our framework encompasses

time-varying network topologies that allow local updates. Therefore, it appears this work would be the first to establish an explicit mean-square deviation (MSD) expression for the federated learning scenario, as well as similar asynchronous learning algorithms that incorporate local updates.

In the upcoming sections, we describe the asynchronous network and illustrate how the original federated learning algorithms initially presented in [30] can be understood as a specific example within this broader framework. Moving forward, the second section provides evidence of the stability of the asynchronous adapt-then-combine (ATC) diffusion algorithm in the federated learning setting. In the subsequent third section, we conduct a comprehensive performance analysis, culminating in the derivation of an MSD expression for the particular case of federated learning. Lastly, in the fourth section, we execute a series of experiments to further explore these concepts.

2. ASYNCHRONOUS NETWORK

2.1. Problem Setup

We consider a network of K agents, similar to the federated or fully decentralized setting, all aiming to solve the convex optimization problem presented as follows:

$$\min_w \frac{1}{K} \sum_{k=1}^K J_k(w), \quad (1)$$

where the local risk function $J_k(\cdot)$ is defined as an empirical average of the local loss function $Q_k(\cdot; x_{k,n})$ over the local dataset $\{x_{k,n}\}_{n=1}^{N_k}$. The communication among agents is restricted by an underlying graph structure, where the combination matrix denoted by A holds elements $a_{\ell k}$ representing the weight agent k assigns to information shared by agent ℓ . We impose certain assumptions on the structure of the underlying graph and the nature of the risk and loss functions.

Assumption 1 (Combination matrix). *The combination matrix is left-stochastic, namely $\mathbb{1}^\top A = \mathbb{1}$ with $a_{k\ell} \geq 0$.* \square

Assumption 2 (Risk and loss functions). *The empirical risks $J_k(\cdot)$ are ν -strongly convex. The loss functions $Q_k(\cdot; \cdot)$ are convex and twice differentiable, namely, for some $\nu > 0$:*

$$J_k(w_2) \geq J_k(w_1) + \nabla_{w^\top} J_k(w_1)(w_2 - w_1) + \frac{\nu}{2} \|w_2 - w_1\|^2, \quad (2)$$

$$Q_k(w_2; \cdot) \geq Q_k(w_1; \cdot) + \nabla_{w^\top} Q_k(w_1; \cdot)(w_2 - w_1). \quad (3)$$

Furthermore, the loss functions have δ -Lipschitz continuous gradients:

$$\|\nabla_{w^\top} Q_k(w_2; x_{k,n}) - \nabla_{w^\top} Q_k(w_1; x_{k,n})\| \leq \delta \|w_2 - w_1\|. \quad (4)$$

\square

School of Engineering, École Polytechnique Fédérale de Lausanne (e-mail: {elsa.rizk, ali.sayed}@epfl.ch).

Center for Machine Learning Research, Peking University (e-mail: kunyuan@pku.edu.cn).

Assumption 3 (Bounded and smooth Hessians). *The Hessians have bounded eigenvalues:*

$$\lambda_{\min} \leq \lambda(\nabla_{w^\top}^2 J_k(w)) \leq \lambda_{\max}, \quad (5)$$

and are locally Lipschitz in a small neighbourhood around w° later defined in (13), namely, there exists $\kappa > 0$ for small Δw :

$$\|\nabla_{w^\top}^2 J_k(w^\circ + \Delta w) - \nabla_{w^\top}^2 J_k(w^\circ)\| \leq \kappa \|\Delta w\|. \quad (6)$$

□




Fig. 1: Illustration of an asynchronous network whose nodes and links change with time.

We introduce the following assumptions on the agents' mode of operation. During an iteration i of the algorithm, an agent k has the option to engage. In the event of participation, it may sample a subset of its neighbourhood \mathcal{N}_k , from which it will aggregate their messages. We attribute a probability q_k to the participation of agent k and we let $q_{\ell k}$ be the sampling probability of agent ℓ by k . Furthermore, prior to any combination step, agent k runs a total of T local update steps. This setup can be modeled as an asynchronous network, where the combination matrix is time varying and random. Accordingly, we let $t \in 1, 2, \dots, T$ denote the local iterations while i denotes the global iterations. The time varying and random combination matrix is written as $\mathbf{A}_{(i-1)T+t}$. Throughout the local iterations, i.e., for $t \neq T$, the combination matrix simplifies to the identity matrix, i.e., $\mathbf{A}_{(i-1)T+t} = \mathbf{I}$. Yet, once $t = T$, a combination step follows after the T th update step. As a result, the combination matrix \mathbf{A}_{iT} will be a sampled version of the original combination matrix \mathbf{A} . For example, in Fig. 1b top left, agent m chooses to participate and it samples agents ℓ and k while leaving out agent 5. Therefore:

$$\begin{aligned} \mathbf{a}_{\ell m, iT} &= a_{\ell m}, & \mathbf{a}_{km, iT} &= a_{km}, & \mathbf{a}_{5m, iT} &= 0, \\ \mathbf{a}_{mm, iT} &= 1 - a_{\ell m} - a_{km}. \end{aligned} \quad (7)$$

Since agent 1 is not participating, then all the weights it attributes to its neighbours will be 0 and its self-weight will be 1, i.e.:

$$\mathbf{a}_{11, iT} = 1, \quad \mathbf{a}_{21, iT} = \mathbf{a}_{\ell 1, iT} = 0. \quad (8)$$

As such, at $t = T$, the elements of the matrix will be given by:

$$\mathbf{a}_{\ell k, iT} = \begin{cases} a_{\ell k}, & \text{with probability } q_k q_{\ell k} \\ 1 - \sum_{m \in \mathcal{N}_k} \mathbf{a}_{mk, iT}, & k = \ell \text{ with probability } q_k \\ 1, & k = \ell \text{ with probability } 1 - q_k \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

To ensure the matrix remains left-stochastic, each agent modifies its self-weight based on the neighbours it has sampled. Moreover, the step-size is also time varying and random:

$$\mu_{k, (i-1)T+t} = \begin{cases} \mu, & \text{with probability } q_k \\ 0, & \text{otherwise} \end{cases} \quad (10)$$

If we consider the ATC diffusion algorithm, the asynchronous version of it can thus be described as follows:

$$\begin{aligned} \psi_{k, (i-1)T+t} &= \mathbf{w}_{k, (i-1)T+t-1} \\ &\quad - \mu_{k, (i-1)T+t} \widehat{\nabla_{w^\top} J_k}(\mathbf{w}_{k, (i-1)T+t-1}), \end{aligned} \quad (11)$$

$$\mathbf{w}_{k, (i-1)T+t} = \sum_{\ell \in \mathcal{N}_k} \mathbf{a}_{\ell k, (i-1)T+t} \psi_{\ell, (i-1)T+t}. \quad (12)$$

Since the combination matrix is left-stochastic, the algorithm does not converge to the minimizer of the average of the risk functions – the solution of problem (1). Instead we can show that on average it will converge to the solution of the weighted average of the risk functions multiplied by the participation probabilities, where the weights are the entries of the Perron eigenvector of the mean combination matrix. Therefore, we let $\bar{p} = \text{col}\{\bar{p}_k\}$ be the Perron eigenvector of $\mathbb{E}\mathbf{A}_{iT}$ and define:

$$w^\circ \triangleq \underset{w}{\text{argmin}} \sum_{k=1}^K \bar{p}_k q_k J_k(w). \quad (13)$$

Federated learning. The following framework can be applied to the federated learning paradigm. The master-slave configuration can be viewed as a fully connected network, where every agent is a neighbour of every other agent. Additionally, the sampling size matches the number of participating agents, signifying that agents are not required to sample their neighbours. By framing the federated system within this interpretation, we can apply the findings discovered in this study.

In the original FedSGD algorithm [30], each agent actively participates in every iteration, executing local update steps. Consequently, all probabilities are uniformly set to 1, $q_k = q_{\ell k} = 1$. As a result the step-size remains constant $\mu_{k, (i-1)T+t} = \mu$. Additionally, the combination matrix alternates between the identity matrix and the full combination matrix with equally weighted entries:

$$\mathbf{A}_{(i-1)T+t} = \begin{cases} \frac{1}{K} \mathbf{1} \mathbf{1}^\top, & t = T \\ \mathbf{I}, & t \neq T. \end{cases} \quad (14)$$

While, the FedAvg algorithm permits agent dropouts, resulting in non-unitary participation probabilities. The sampling probabilities $q_{\ell k}$ are set to 1. Consequently, during the global iteration i , if L_i agents participate, the combination weights at $t = T$ are given by:

$$\mathbf{a}_{\ell k, iT} = \begin{cases} \frac{1}{L_i}, & \ell \text{ participating} \\ 1, & k = \ell \text{ not participating} \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

For $t \neq T$, the combination matrix remains identity. The step-size operates similarly as in (10). Thus in this particular setting, the combination matrix remains doubly-stochastic during each iteration. As such, the entries of the Perron eigenvector \bar{p}_k of the mean combination matrix are equal to $1/K$, and the optimal model w° simplifies to the solution of the original optimization problem (1).

2.2. Stability Analysis

We begin by establishing the stability of the algorithm in the mean-square error sense. This involves starting with the formulation of the error recursion and then proceeding to define the gradient noise. Consequently, we define the error as $\tilde{\mathbf{w}}_{k,(i-1)T+t} = w^o - \mathbf{w}_{k,(i-1)T+t}$ and the gradient noise as:

$$\mathbf{s}_{k,(i-1)T+t} \triangleq \widehat{\nabla_{w^\top} J_k}(\mathbf{w}_{k,(i-1)T+t-1}) - \nabla_{w^\top} J_k(\mathbf{w}_{k,(i-1)T+t-1}). \quad (16)$$

Then, by invoking the mean-value theorem [31], we can express the gradient as:

$$\nabla_{w^\top} J_k(\mathbf{w}_{k,(i-1)T+t-1}) = -\mathbf{H}_{k,(i-1)T+t-1} \tilde{\mathbf{w}}_{k,(i-1)T+t-1} - \nabla_{w^\top} J_k(w^o), \quad (17)$$

where we define the following terms:

$$\overline{\mathbf{H}}_{k,(i-1)T+t} \triangleq (I - \boldsymbol{\mu}_{k,(i-1)T+t} \mathbf{H}_{k,(i-1)T+t-1}), \quad (18)$$

$$\mathbf{H}_{k,(i-1)T+t-1} \triangleq \int_0^1 \nabla_{w^\top}^2 J_k(w^o - \tau \tilde{\mathbf{w}}_{k,(i-1)T+t-1}) d\tau. \quad (19)$$

Accordingly, the expression of the error recursion could be formulated as:

$$\begin{aligned} \tilde{\mathbf{w}}_{k,(i-1)T+t} &= \sum_{\ell \in \mathcal{N}_k} \mathbf{a}_{\ell k,(i-1)T+t} \left(\overline{\mathbf{H}}_{\ell,(i-1)T+t} \tilde{\mathbf{w}}_{\ell,(i-1)T+t-1} \right. \\ &\quad \left. + \boldsymbol{\mu}_{\ell,(i-1)T+t} (\mathbf{s}_{\ell,(i-1)T+t} - \nabla_{w^\top} J_\ell(w^o)) \right). \end{aligned} \quad (20)$$

Initially, we demonstrate that the stochastic gradient is an unbiased estimate of the true gradient and that the gradient noise has a finite second-order moment.

Lemma 1 (First and second-order moments of gradient noise). *The gradient noise defined in (16) has zero-mean and bounded second-order moment, namely:*

$$\mathbb{E} \|\mathbf{s}_{k,(i-1)T+t}\|^2 \leq \beta_s^2 \mathbb{E} \|\tilde{\mathbf{w}}_{k,(i-1)T+t-1}\|^2 + \sigma_s^2, \quad (21)$$

where β_s^2 and σ_s^2 are some constants.

Proof. Proof omitted due to space limitations. \square

Following arguments similar to [31], it can be demonstrated that the algorithm achieves exponential convergence to a region around the true model.

Theorem 1 (Mean-square stability). *Under assumptions 1, 2, 3, and for small enough step-size:*

$$\mu \leq \frac{2\lambda_{\min}}{\lambda_{\max}^2 + \beta_s^2}, \quad (22)$$

the individual errors converge exponentially fast:

$$\limsup_{i \rightarrow \infty} \mathbb{E} \|\tilde{\mathbf{w}}_{k,(i-1)T+t}\|^2 \leq \frac{\sigma_s^2}{1-\gamma} \mu^2 = O(\mu), \quad (23)$$

where $\gamma \triangleq \max_k 1 - 2\mu q_k \lambda_{\min} + \mu^2 q_k (\lambda_{\max}^2 + \beta_s^2) \in [0, 1)$ is the convergence rate.

Proof. Proof omitted due to space limitations. \square

As is evident from the theorem's statement, the rate and region of convergence are determined by the least active agent. To put it differently, the agent with the lowest participation probability q_k slows down the overall algorithm. Moreover, given this agent's infrequent participation, it negatively impacts the overall network performance. Consequently, the strength of the network is contingent on its most fragile component.

In order to conduct the performance analysis, it is necessary to examine the fourth-order stability of the algorithm. Thus, by assuming that the gradient noise has bounded fourth-order moment, we can further expand upon the previous result in a manner similar to [31].

Theorem 2 (Fourth-order stability). *If the gradient noise has a bounded fourth-order moment:*

$$\mathbb{E} \|\mathbf{s}_{k,(i-1)T+t}\|^4 \leq \beta_s^4 \mathbb{E} \|\tilde{\mathbf{w}}_{k,(i-1)T+t-1}\|^4 + \sigma_s^4, \quad (24)$$

and for small enough step-size, then:

$$\limsup_{i \rightarrow \infty} \mathbb{E} \|\tilde{\mathbf{w}}_{k,(i-1)T+t}\|^4 \leq O(\mu^2). \quad (25)$$

Proof. Proof omitted due to space limitations. \square

2.3. Performance Analysis

We are now ready to proceed with the performance analysis, wherein we will present an expression for the MSD for the federated setting. However, before delving into this, we lay out the following assumption concerning the noise process.

Assumption 4 (Noise process). *Define the covariance of the gradient noise:*

$$R_{k,(i-1)T+t}(w) \triangleq \mathbb{E} \mathbf{s}_{k,(i-1)T+t}(w) \mathbf{s}_{k,(i-1)T+t}^\top(w). \quad (26)$$

Then, for some positive constants κ_s and α_s , the covariance satisfies the following Lipschitz condition:

$$\begin{aligned} &\|\text{diag}\{R_{k,(i-1)T+t}(w^o) - R_{k,(i-1)T+t}(\mathbf{w}_{k,(i-1)T+t-1})\}\| \\ &\leq \kappa_s \|\text{col}\{\tilde{\mathbf{w}}_{k,(i-1)T+t-1}\}\|^{\alpha_s}, \end{aligned} \quad (27)$$

and the following limit exists:

$$R_k \triangleq \lim_{i \rightarrow \infty} R_{k,(i-1)T+t}(w^o). \quad (28)$$

\square

Using the aforementioned assumption and the smoothness assumption of the Hessians, we can derive an expression for the MSD. We introduce the matrix \mathcal{G}_t :

$$\begin{aligned} \mathcal{G}_t &\triangleq \mathbb{E} \left\{ \mathcal{A}_{(i-1)T+t}^\top (I - \mathcal{M}_{k,(i-1)T+t} \text{diag}\{\nabla_{w^\top}^2 J_k(w^o)\}) \right. \\ &\quad \left. \otimes_b \mathcal{A}_{(i-1)T+t}^\top (I - \mathcal{M}_{k,(i-1)T+t} \text{diag}\{\nabla_{w^\top}^2 J_k(w^o)\}) \right\}, \end{aligned} \quad (29)$$

where the operator \otimes_b represents the block Kronecker product and:

$$\mathcal{A}_{(i-1)T+t} \triangleq \mathbf{A}_{(i-1)T+t} \otimes I, \quad (30)$$

$$\mathcal{M}_{(i-1)T+t} \triangleq \text{diag}\{\boldsymbol{\mu}_{k,(i-1)T+t}\}. \quad (31)$$

For $t \neq T$, the matrix simplifies to:

$$\mathcal{G}_t = \text{diag} \{G_{\ell k}\}, \quad (32)$$

$$G_{\ell k} \triangleq \begin{cases} (I - \mu q_\ell \nabla_{w^\top}^2 J_\ell(w^o))(I - \mu q_k \nabla_{w^\top}^2 J_k(w^o)), & \ell \neq k \\ I - q_k \mu \nabla_{w^\top}^2 J_k(w^o), & \ell = k \end{cases} \quad (33)$$

At $t = T$, the matrix \mathcal{G}_T captures the interdependencies among neighbours and the impact of neighbour sampling probabilities $q_{\ell k}$. Due to the complexity of the expression of \mathcal{G}_T , we omit its inclusion.

We next introduce the matrix \mathcal{C}_t which captures the dependency between the step-size and combination weights:

$$\mathcal{C}_t \triangleq \mathbb{E} \{ (\mathcal{A}_{(i-1)T+t} \otimes_b \mathcal{A}_{(i-1)T+t}) \times (\mathcal{M}_{(i-1)T+t} \otimes_b \mathcal{M}_{(i-1)T+t}) \}. \quad (34)$$

For similar reasons, we refrain from explicitly formulating the matrix expression at $t = T$. Nonetheless, for $t \neq T$, we have:

$$\mathcal{C}_t = \text{diag} \{C_{\ell k}\}, \quad (35)$$

$$C_{\ell k} \triangleq \begin{cases} q_k q_\ell \mu^2, & \ell \neq k \\ q_k \mu^2, & \ell = k \end{cases} \quad (36)$$

Subsequently, we can formulate an expression for the MSD.

Theorem 3 (Steady-state MSD). *It holds that:*

$$\text{MSD} = \frac{1}{K} z^\top \text{bvec}(I) + O(\mu^{1+\alpha_0}), \quad (37)$$

where the block vectorization operator bvec stacks the columns of the blocks of the matrix, and:

$$\alpha_0 \triangleq \frac{1}{2} \min\{1, \alpha_s\}, \quad (38)$$

$$z \triangleq \left(I - \left(\mathcal{G}_T \mathcal{G}_t^{T-1} \right)^\top \right)^{-1} \mathcal{C} \text{bvec}(\text{diag}\{R_k\}), \quad (39)$$

$$\mathcal{C} \triangleq \left(I + \left(\mathcal{G}_T \mathcal{G}_t^{T-1} \right)^\top \right) \mathcal{C}_T + \mathcal{G}_T^\top \sum_{j=1}^{T-1} \left(\mathcal{G}_t^{j-1} \right)^\top \mathcal{C}_t. \quad (40)$$

Proof. Proof omitted due to space limitations. \square

3. EXPERIMENTAL RESULTS

We consider a linear regression problem of the form:

$$\min_w \frac{1}{K} \sum_{k=1}^K \frac{1}{N} \sum_{n=1}^N \|\mathbf{d}_k(n) - \mathbf{u}_{k,n}^\top w\|^2. \quad (41)$$

We generate for each agent a data set $\{\mathbf{d}_k(n), \mathbf{u}_{k,n}\}$ consisting of $N = 10^6$ samples. These samples include five dimensional feature vectors $\mathbf{u}_{k,n}$ drawn from a normal distribution $\mathcal{N}(0, R_u)$ and an independent Gaussian noise $\mathbf{v}_k(n) \sim \mathcal{N}(0, \sigma_{v,k}^2)$. A random generative model w^* is sampled from $\mathcal{N}(0, R_w)$ and the labels are determined by $\mathbf{d}_k(n) = \mathbf{u}_{k,n}^\top w^* + \mathbf{v}_k(n)$. For a linear regression problem, the optimal model w^o can be calculated using \hat{R}_u and \hat{r}_{uv} , which represent the sample covariance and cross-covariance:

$$w^o = w^* + \hat{R}_u^{-1} \hat{r}_{uv}. \quad (42)$$




Fig. 2: MSD curve for the asynchronous ATC diffusion algorithm.

The network comprises of $K = 20$ agents. We consider three different cases of the asynchronous network. Case 1 is the most general case where agent subsampling and local updates occur. The participation probabilities are set to $q_k = 0.5$. Neighbour sampling probabilities $q_{\ell k}$ are randomly assigned. Finally, the parameter T is chosen as 100. Case 2 only considers agent subsampling and no local updates. Thus, $T = 1$ and the probabilities are kept as in case 1. Finally, case 3 assumes full agent participation with no subsampling of neighbourhoods ($q_k = q_{\ell k} = 1$), while allows local updates ($T = 100$). This case coincides with the FedSGD algorithm.

Running the algorithm using a step-size $\mu = 0.0001$, we calculate the resultant MSD throughout the algorithm's progression. Afterward, we average the MSD across 5 experiments and illustrate the average curve alongside the theoretical MSD expression found in Theorem 3. Investigating Fig. 2, we observe that as time passes the approximated MSD approaches the theoretical value, even for the general decentralized case. Moreover, in the context of the given setup, whether local updates are used, as in case 1, or not, as in case 2, an identical behaviour is exhibited. Nonetheless, the primary factor affecting speedup becomes the agents' participation, particularly evident in case 3, where all agents participate. Consequently, the adoption of local updates has no discernible impact on the convergence rate of the algorithm. Furthermore, the three algorithms observe comparable theoretical MSD due to the small employed step-size.

4. CONCLUSION

In summary, this study focuses on asynchronous networks, where the primary goal is to solve a learning problem framed as an optimization task. Our investigation revolves around the premise that each agent autonomously decides when to participate in the algorithm and with which subset of its neighbourhood it will collaborate. Additionally, agents perform local updates before sharing their results. In this framework, we observe that in the federated setting the algorithm remains stable by continuously converging to an $O(\mu)$ neighbourhood of the optimal model w^o , just like the synchronous version of the algorithm. However, the rate and region of convergence are now influenced by the frequency of the nonparticipation of agents as well as the degree of connectedness of the network. Furthermore, these effects are aggregated in an actual expression of the MSD.

5. REFERENCES

- [1] D. Bertsekas, "A new class of incremental gradient methods for least squares problems," *SIAM J. Optim.*, vol. 7, pp. 913–926, Nov. 1997.
- [2] F. S. Cattivelli and A. H. Sayed, "Analysis of spatial and incremental LMS processing for distributed estimation," *IEEE Transactions on Signal Processing*, vol. 59, no. 4, pp. 1465–1480, 2011.
- [3] C. G. Lopes and A. H. Sayed, "Incremental adaptive strategies over distributed networks," *IEEE Transactions on Signal Processing*, vol. 55, no. 8, pp. 4064–4077, 2007.
- [4] E. S. Helou and A. R. De Pierro, "Incremental subgradients for constrained convex optimization: A unified framework and new methods," *SIAM J. Optim.*, vol. 20, no. 3, p. 1547–1572, Dec 2009.
- [5] A. Nedic and D. P. Bertsekas, "Incremental subgradient methods for nondifferentiable optimization," *SIAM Journal on Optimization*, vol. 12, no. 1, pp. 109–138, 2001.
- [6] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association.*, vol. 69, no. 345, pp. 118–121, 1974.
- [7] A. Nedic and A. Ozdaglar, "Distributed subgradient methods for multi-agent optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, 2009.
- [8] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Systems & Control Letters*, vol. 53, no. 1, pp. 65–78, Sep 2004.
- [9] B. Johansson, T. Keviczky, M. Johansson, and K. H. Johansson, "Subgradient methods and consensus algorithms for solving convex optimization problems," in *Proc. IEEE Conf. Dec. Control (CDC)*, Cancun, Mexico, December 2008, pp. 4185–4190.
- [10] C. G. Lopes and A. H. Sayed, "Diffusion least-mean squares over adaptive networks: Formulation and performance analysis," *IEEE Transactions on Signal Processing*, vol. 56, no. 7, pp. 3122–3136, Jul 2008.
- [11] J. Chen and A. H. Sayed, "Diffusion adaptation strategies for distributed optimization and learning over networks," *IEEE Transactions on Signal Processing*, vol. 60, no. 8, pp. 4289–4305, Aug 2012.
- [12] S.-Y. Tu and A. H. Sayed, "Diffusion strategies outperform consensus strategies for distributed estimation over adaptive networks," *IEEE Transactions on Signal Processing*, vol. 60, no. 12, pp. 6217–6234, Dec 2012.
- [13] J. Chen and A. H. Sayed, "On the learning behavior of adaptive networks—part i: Transient analysis," *IEEE Transactions on Information Theory*, vol. 61, no. 6, pp. 3487–3517, Dec 2015.
- [14] S. Vlaski and A. H. Sayed, "Distributed learning in non-convex environments—part i: Agreement at a linear rate," *IEEE Transactions on Signal Processing*, vol. 69, pp. 1242–1256, 2021.
- [15] K. I. Tsianos, S. Lawlor, and M. G. Rabbat, "Consensus-based distributed optimization: Practical issues and applications in large-scale machine learning," in *Annual Allerton Conference on Communication, Control, and Computing*, 2012, pp. 1543–1550.
- [16] C. N. Hadjicostis and T. Charalambous, "Average consensus in the presence of delays in directed graph topologies," *IEEE Transactions on Automatic Control*, vol. 59, no. 3, pp. 763–768, 2013.
- [17] K. Tsianos, S. Lawlor, and M. Rabbat, "Communication/computation tradeoffs in consensus-based distributed optimization," *Advances in Neural Information Processing Systems*, vol. 25, pp. 1943–1951, 2012.
- [18] M. Carpentiero, V. Matta, and A. H. Sayed, "Compressed distributed regression over adaptive networks," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, Rhodes, Greece, 2023, pp. 1–5.
- [19] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Transactions on Automatic Control*, vol. 31, no. 9, pp. 803–812, 1986.
- [20] K. Srivastava and A. Nedic, "Distributed asynchronous constrained stochastic optimization," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 4, pp. 772–790, Aug 2011.
- [21] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise," *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 355–369, 2009.
- [22] S. Kar and J. M. Moura, "Distributed consensus algorithms in sensor networks: Quantized data and random link failures," *IEEE Transactions on Signal Processing*, vol. 58, no. 3, pp. 1383–1400, 2009.
- [23] R. Z. Aviv, I. Hakimi, A. Schuster, and K. Y. Levy, "Asynchronous distributed learning : Adapting to gradient delays without prior knowledge," in *Proceedings of the International Conference on Machine Learning*, July 2021, pp. 436–445.
- [24] M. S. Assran and M. G. Rabbat, "Asynchronous gradient push," *IEEE Transactions on Automatic Control*, vol. 66, no. 1, pp. 168–183, 2021.
- [25] X. Lian, W. Zhang, C. Zhang, and J. Liu, "Asynchronous decentralized parallel stochastic gradient descent," in *Proceedings of the International Conference on Machine Learning*, vol. 80, Stockholm, Sweden, July 2018, pp. 3043–3052.
- [26] X. Zhao and A. H. Sayed, "Asynchronous adaptation and learning over networks—part i: Modeling and stability analysis," *IEEE Transactions on Signal Processing*, vol. 63, no. 4, pp. 811–826, 2014.
- [27] C. G. Lopes and A. H. Sayed, "Randomized incremental protocols over adaptive networks," in *Proc. IEEE ICASSP*, Dallas, TX, March 2010, pp. 3514–3517.
- [28] Ø. L. Rørtveit, J. H. Husøy, and A. H. Sayed, "Diffusion LMS with communication constraints," in *Proc. 44th Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, Nov. 2010, pp. 1645–1649.
- [29] R. Arablouei, S. Werner, K. Doğançay, and Y.-F. Huang, "Analysis of a reduced-communication diffusion LMS algorithm," *Signal Processing*, vol. 117, pp. 355–361, 2015.
- [30] H. B. McMahan, E. Moore, D. Ramage, and S. Hampson, "Communication-efficient learning of deep networks from decentralized data," *Proc. International Conference on Artificial Intelligence and Statistics*, vol. 54, pp. 1273–1282, 20–22 April 2017.
- [31] A. H. Sayed, "Adaptation, learning, and optimization over networks," *Foundations and Trends in Machine Learning*, vol. 7, no. 4-5, pp. 311–801, 2014.