

Composite Optimization in Open Multi-Agent Systems under Additive Errors

Diego Deplano, *Member, IEEE*, Nicola Bastianello, *Member, IEEE*
Mauro Franceschelli, *Senior, IEEE*, and Karl H. Johansson, *Fellow, IEEE*

Abstract—This paper studies distributed composite optimization in open multi-agent systems, where agents may join and leave the network over time. We propose a variant of Open ADMM in which the composite proximal step is approximated through a finite number of proximal-gradient iterations, enabling agents to handle composite objectives even when the exact proximal operator is not available in closed form. Since the inexact computation of the proximal step induces additive errors, we establish a convergence result for paracontractive open multi-agent systems subject to bounded additive perturbations that capture both inexact local updates and noisy communications. Under mild assumptions on the network topology, the proposed algorithm converges linearly to a neighborhood of the optimal consensus set, with an explicit bound on the steady-state tracking error. The theoretical guarantees are corroborated by numerical experiments on distributed learning applications, namely linear regression (supervised learning) and principal component analysis (unsupervised learning).

Index Terms—Robust optimization, distributed optimization, online optimization, online learning, classification, PCA, open networks.

I. INTRODUCTION

Distributed optimization over networks is a core primitive in large-scale learning and control, enabling agents to cooperatively estimate a common decision variable using only local objectives and peer-to-peer communication. In many learning and estimation tasks, the local costs naturally decompose into a *composite* form, consisting of a smooth loss term plus a (possibly non-smooth) regularizer or constraint, typically modeled by a *proximable* function [1], [2], [3], [4], [5], [6], [7], [8], [9]. This motivates the study of distributed *composite* optimization problems over networks of interacting agents, who seek agreement on a common model while leveraging local data and local structure.

The work of Diego Deplano and Mauro Franceschelli was partially funded by Fondazione di Sardegna with project “Formal methods for Resiliency and Security in Networks of Systems (ReSNetS)” call 2025, F83C26000350007.

The work of Nicola Bastianello and Karl H. Johansson was supported by the Swedish Research Council Distinguished Professor Grant 2017-01078 Knut and Alice Wallenberg Foundation Wallenberg Scholar Grant.

Diego Deplano and Mauro Franceschelli are with DIEE, University of Cagliari, 09123 Cagliari, Italy. Emails: {diego.deplano,mauro.franceschelli}@unica.it

Nicola Bastianello and Karl H. Johansson are with the School of Electrical Engineering and Computer Science, and Digital Futures, KTH Royal Institute of Technology, Stockholm, Sweden. Emails: {nicolba,kallej}@kth.se

The current literature on distributed composite optimization is limited to *closed* networks with a fixed set of agents, whereas practical deployments are often *open*: agents may join or leave due to mobility, availability, failures, privacy policies, economic incentives, and so on. Such scenarios fall under the umbrella of *open multi-agent systems*, where the number of participating agents is time-varying and classical convergence guarantees do not apply. In open networks, the consensus optimizer becomes inherently time-dependent, and join/leave events can prevent convergence, generate persistent transients, and even destabilize algorithms that are convergent in the static setting. Recent work has begun to formalize stability and performance notions tailored to open networks, and to develop analysis tools that expose the arising challenges [10], [11], [12], [13], while most contributions in distributed optimization for open networks remain algorithm- and application-specific; a broad account of these developments, which address non-composite optimization or consensus, is provided in [13].

The **main contribution** of this paper is to extend and analyze Open ADMM for distributed composite optimization problems over open networks. The specific novel contributions of the paper are as follows:

- 1) we specialize Open ADMM presented by us in [13] to deal with composite problems by approximating the local proximal update via a finite number of proximal-gradient inner iterations (Algorithm 1);
- 2) we characterize Open ADMM’s convergence and performance by deriving an explicit bound that links tracking accuracy to the variability of the optimal solutions and to disagreement between local and global minimizers, and by showing linear convergence toward a neighborhood of the optimal consensus set (Theorem 1);
- 3) to prove the above, we establish a convergence result for paracontractive open multi-agent dynamics subject to bounded additive perturbations (Theorem 2), which constitutes a general analysis and design tool;
- 4) we showcase Open ADMM’s performance on two representative learning tasks with local constraints: supervised linear regression and unsupervised distributed principal component analysis.

II. PRELIMINARIES ON OPEN MULTI-AGENT SYSTEMS

The sets of real and natural numbers are denoted by $\mathbb{R}$ and $\mathbb{N}$, respectively. Their positive subsets are denoted by $\mathbb{R}_+$ and $\mathbb{N}_+$, respectively. For examples and visual illustrations of the preliminaries reported next, we refer to [13, Sec. II].

Let $\mathcal{I} \neq \emptyset$ be a finite label set. Then, $x \in \mathbb{R}^{\mathcal{I}}$ denotes a point with labeled components $x^i, i \in \mathcal{I}$. To compare points in different labeled spaces, we use the concept of *open distance* from [13, Definition 3].

Definition 1 (Open distance). The open distance between $x \in \mathbb{R}^{\mathcal{I}_1}, y \in \mathbb{R}^{\mathcal{I}_2}$, with $\mathcal{I}_1 \cap \mathcal{I}_2 \neq \emptyset$, is $d(x, y) = \|\tilde{x} - \tilde{y}\|_2$ where $\tilde{x} = [x^i | i \in \mathcal{I}_1 \cap \mathcal{I}_2]$ and $\tilde{y} = [y^i | i \in \mathcal{I}_1 \cap \mathcal{I}_2]$.

Given $x \in \mathbb{R}^{\mathcal{I}_0}, \mathcal{X} \subseteq \mathbb{R}^{\mathcal{I}_1}, \mathcal{Y} \subseteq \mathbb{R}^{\mathcal{I}_2}$, the distance from a point to a set and the distance between two sets, of possibly different labeled spaces, are defined by $d(x, \mathcal{Y}) := \inf_{y \in \mathcal{Y}} d(x, y)$ and $d(\mathcal{X}, \mathcal{Y}) := \inf_{x \in \mathcal{X}} d(x, \mathcal{Y})$. The set of points in $\mathcal{Y}$ at minimum distance from x is called the projection of x onto $\mathcal{Y}$ and is given by the *projection operator* $\text{proj}(x, \mathcal{Y}) = \{y^* \in \mathcal{Y} : d(x, y^*) = \inf_{y \in \mathcal{Y}} d(x, y)\}$. The projection operator enables the definition of the *shadow distance* between sets [13, Definition 4], capturing the worst-case distance between the corresponding projection points.

Definition 2 (Shadow distance). The shadow distance between sets $\mathcal{X} \subseteq \mathbb{R}^{\mathcal{I}_1}, \mathcal{Y} \subseteq \mathbb{R}^{\mathcal{I}_2}$ with $\mathcal{I}_1 \cap \mathcal{I}_2 \neq \emptyset$ is

$$d_{\text{sh}}(\mathcal{X}, \mathcal{Y}) := \sup_{z \in \mathbb{R}^{\mathcal{I}_1 \cup \mathcal{I}_2}} d(\text{proj}(z, \mathcal{X}), \text{proj}(z, \mathcal{Y})).$$

The shadow distance yields the following triangle inequality:

$$d(z, \mathcal{X}) \leq d(z, \mathcal{Y}) + d_{\text{sh}}(\mathcal{X}, \mathcal{Y}). \quad (1)$$

The projection operator is a special case of the *proximal operator*, which, for a point $x \in \mathbb{R}^{\mathcal{I}_1}$, a function $f : \mathcal{Y} \rightarrow \mathbb{R}$ with $\mathcal{Y} \subseteq \mathbb{R}^{\mathcal{I}_2}$, and a penalty parameter $\varepsilon > 0$, is defined by

$$\text{prox}_f^\varepsilon(x) = \left\{ y^* \in \mathcal{Y} : f(y^*) + \frac{d^2(x, y^*)}{2\varepsilon} = \min_{y \in \mathcal{Y}} \left(f(y) + \frac{d^2(x, y)}{2\varepsilon} \right) \right\}. \quad (2)$$

Indeed, $\text{proj}(x, \mathcal{Y}) = \text{prox}_{\iota_{\mathcal{Y}}}^\varepsilon(x)$ for $\iota_{\mathcal{Y}} : \mathbb{R}^{\mathcal{I}_2} \rightarrow \mathbb{R} \cup \{+\infty\}$ defined as $\iota_{\mathcal{Y}}(y) = 0$ if $y \in \mathcal{Y}$, and $\iota_{\mathcal{Y}}(y) = +\infty$ otherwise.

An *open multi-agent system* (OMAS) is a networked system with a time-varying number $n_k \in \mathbb{N}_+$ of agents over time $k \in \mathbb{N}$. Interactions at time k are modeled by an undirected time-varying graph $\mathcal{G}_k = (\mathcal{V}_k, \mathcal{E}_k)$, where $\mathcal{V}_k$ is a finite node set with $|\mathcal{V}_k| = n_k$ and $\mathcal{E}_k \subseteq \mathcal{V}_k \times \mathcal{V}_k$ is the edge set. The neighbor set and degree of agent i are $\mathcal{N}_k^i := \{j \in \mathcal{V}_k \setminus \{i\} : (i, j) \in \mathcal{E}_k\}$ and $\eta_k^i := |\mathcal{N}_k^i|$. The agent set is partitioned into *remaining* $\mathcal{R}_k := \mathcal{V}_k \cap \mathcal{V}_{k-1}$, *arriving* $\mathcal{A}_k := \mathcal{V}_k \setminus \mathcal{V}_{k-1}$, and *departing* $\mathcal{D}_k := \mathcal{V}_k \setminus \mathcal{V}_{k+1}$ agents.

Each agent $i \in \mathcal{V}_k$ has a state $x_k^i \in \mathbb{R}^{\mathcal{I}_k^i}$ of time-varying dimension $|\mathcal{I}_k^i|$, where $\mathcal{I}_k^i$ is the label set of its components. Without loss of generality, we let $\mathcal{I}_k^i = \mathcal{S}_k^i \times \mathcal{P}$, where $\mathcal{S}_k^i$ is a finite time-varying label set indexing the state blocks of agent i , and $\mathcal{P}$ is a fixed finite label set indexing the components of each block, so that $x_k^i = [\dots, x_k^{ij}, \dots]$ stacks vectors $x_k^{ij} \in \mathbb{R}^{\mathcal{P}}$ for all $i \in \mathcal{V}_k, j \in \mathcal{S}_k^i$. For each $i \in \mathcal{V}_k$, the set $\mathcal{S}_k^i$ is partitioned into *remaining* $\mathcal{R}_k^i := \mathcal{S}_k^i \cap \mathcal{S}_{k-1}^i$, *arriving* $\mathcal{A}_k^i := \mathcal{S}_k^i \setminus \mathcal{S}_{k-1}^i$, and *departing* $\mathcal{D}_k^i := \mathcal{S}_k^i \setminus \mathcal{S}_{k+1}^i$ labels of the i -th agent.

The label set indexing the network state $x_k \in \mathbb{R}^{\mathcal{I}_k}$ is $\mathcal{I}_k := \{(i, j, p) \mid i \in \mathcal{V}_k, j \in \mathcal{S}_k^i, p \in \mathcal{P}\}$, and $x_k = [\dots, x_k^i, \dots]$ stacks vectors $x_k^i \in \mathbb{R}^{\mathcal{S}_k^i \times \mathcal{P}}$ for all $i \in \mathcal{V}_k$. The evolution of each network state component $x_k^{ij} \in \mathbb{R}^{\mathcal{P}}$ combines a causal update for remaining labels of remaining agents with an initialization rule for arriving agents and arriving labels of remaining agents,

$$x_k^{ij} = \mathsf{T}_k^{ij}(x_{k-1}) := \begin{cases} \mathsf{F}_k^{ij}(x_{k-1}) & \text{if } i \in \mathcal{R}_k \wedge j \in \mathcal{R}_k^i, \\ x_k^{A, ij} & \text{if } i \in \mathcal{A}_k \vee j \in \mathcal{A}_k^i, \end{cases} \quad (3)$$

which in compact form becomes

$$x_k = \mathsf{T}_k(x_{k-1}), \quad x_k \in \mathbb{R}^{\mathcal{I}_k}. \quad (4)$$

If the agent and label sets are unchanged, then no initialization is needed and the state update rule T_k reduces to the map $\mathsf{F}_k : \mathbb{R}^{\mathcal{I}_{k-1}} \rightarrow \mathbb{R}^{\mathcal{I}_{k-1}}$, yielding the *standard iteration*

$$x_k = \mathsf{F}_k(x_{k-1}), \quad \text{when } \mathcal{I}_k = \mathcal{I}_{k-1}. \quad (5)$$

In open systems, the concept of equilibrium point is replaced by that of a *point of interest*, i.e., a fixed point of the operator F_k in eq. (5). Under the assumption that at each time $k \in \mathbb{N}$ there is only one point of interest, the concept of a *trajectory of points of interest* was introduced in [10, Definition 3.1]; following [13, Definition 1], it can be generalized to the case of multiple points (sets) of interest.

Definition 3 (Trajectory of sets of interest).

Consider an OMAS and assume that for each $k \in \mathbb{N}$ the equation $x = \mathsf{F}_{k+1}(x)$ admits at least one solution $\hat{x}_k \in \mathbb{R}^{\mathcal{I}_k}$, which is called *point of interest*, and let $\hat{\mathcal{X}}_k \subseteq \mathbb{R}^{\mathcal{I}_k}$ be the set of all points of interest at $k \in \mathbb{N}$, i.e.,

$$\hat{\mathcal{X}}_k := \{x \in \mathbb{R}^{\mathcal{I}_k} \mid x = \mathsf{F}_{k+1}(x)\}. \quad (6)$$

Then, the sequence $\{\hat{\mathcal{X}}_k : k \in \mathbb{N}\}$ is called the *trajectory of sets of interest (TSI)* of the OMAS.

The existence of a TSI is guaranteed for paracontractive OMASs, i.e., OMASs whose standard iteration (5) is paracontractive [13, Definition 5].

Definition 4 (Paracontractivity). An OMAS is said to be *paracontractive* if there exists $\gamma \in (0, 1)$ such that for all $k \geq 0$ and for all $x \in \mathbb{R}^{\mathcal{I}_k}$ it holds

$$d(\mathsf{F}_{k+1}(x), \hat{\mathcal{X}}_k) \leq \gamma \cdot d(x, \hat{\mathcal{X}}_k), \quad (7)$$

where $\{\hat{\mathcal{X}}_k : k \in \mathbb{N}\}$ is the TSI of the OMAS.

Finally, stability is tailored to size-varying settings by referring it to the TSI rather than the set of equilibrium points. The following definitions are straightforward generalizations of [10, Definition 3.3] and [14, Definition 5] for OMASs with multiple points (sets) of interest at each time $k \in \mathbb{N}$.

Definition 5 (Open stability). Consider an OMAS with state evolution $\{x_k : k \in \mathbb{N}\}$. Its TSI $\{\hat{\mathcal{X}}_k : k \in \mathbb{N}\}$ is said to be *open stable* if there is a stability radius $R \geq 0$ with the following property: for every $\varepsilon > R$, there is $\delta > 0$ such that:

$$\frac{d(x_0, \hat{\mathcal{X}}_0)}{\sqrt{|\mathcal{I}_0|}} < \delta \Rightarrow \frac{d(x_k, \hat{\mathcal{X}}_k)}{\sqrt{|\mathcal{I}_k|}} < \varepsilon, \quad \forall k \geq 0.$$

Definition 6 (Global asymptotic open stability).

Consider an OMAS whose TSI $\{\hat{\mathcal{X}}_k : k \in \mathbb{N}\}$ is open stable with stability radius $R \geq 0$. The TSI is said to be *globally asymptotically open stable* if all trajectories converge to within a distance of R from the TSI:

$$\limsup_{k \rightarrow \infty} \frac{d(x_k, \hat{\mathcal{X}}_k)}{\sqrt{|\mathcal{I}_k|}} \leq R.$$

III. PROBLEM FORMULATION AND MAIN RESULT

We consider the distributed optimization problem

$$\min_{y_i \in \mathbb{R}^{\mathcal{P}}} \sum_{i \in \mathcal{V}_k} (f_k^i(y_i) + g_k^i(y_i)) \quad \text{s.t. } y_i = y_j \quad \forall (i, j) \in \mathcal{E}_k, \quad (8)$$

where $|\mathcal{P}|$ is the number of variables and $f_k^i + g_k^i : \mathbb{R}^{\mathcal{P}} \mapsto \mathbb{R}$ denotes the local, composite objective function of agent $i \in \mathcal{V}_k$ at time $k \in \mathbb{N}$. We aim to find an OMAS to solve (8) under suitable assumptions on the local objective functions and the corresponding local and global solutions. We formalize our set of assumptions in Assumption 1, which makes use of the following notation for the set of solutions to (8),

$$\mathcal{Y}_k^* := \operatorname{argmin}_{y \in \mathbb{R}^{\mathcal{P}}} \sum_{i \in \mathcal{V}_k} (f_k^i(y) + g_k^i(y)),$$

and for the minimizers of each local cost function,

$$\mathcal{Y}_k^{i,*} = \operatorname{argmin}_{y \in \mathbb{R}^{\mathcal{P}}} (f_k^i(y) + g_k^i(y)). \quad (9)$$

Assumption 1. Problem (8) is such that, $\forall k \in \mathbb{N}$: (i) f_k^i, g_k^i are proper, lower semi-continuous, and convex¹ for all $i \in \mathcal{V}_k$; (ii) f_k^i are uniformly L -smooth over $i \in \mathcal{V}_k$ and $k \in \mathbb{N}$, while functions g_k^i are proximable²; (iii) the set of minimizers $\mathcal{Y}_k^{i,*} \subseteq \mathbb{R}^{\mathcal{P}}$ for each local composite cost function $f_k^i + g_k^i$ is nonempty; (iv) the distance between any two consecutive global solutions $y_k^* \in \mathcal{Y}_k^*$ and $y_{k-1}^* \in \mathcal{Y}_{k-1}^*$ is upper bounded by $\sigma \geq 0$; (v) the distance between any local solution $y_k^{i,*} \in \mathcal{Y}_k^{i,*}$ and any global solution $y_k^* \in \mathcal{Y}_k^*$ is upper bounded by $\omega \geq 0$.

We detail in Algorithm 1 how to apply Open ADMM, presented by us in [13], to solve the composite problem (8). At each time $k \in \mathbb{N}$, each agent $i \in \mathcal{V}_k$ maintains and updates one edge variable $x_k^{ij} \in \mathbb{R}^{\mathcal{P}}$ for each neighbor $j \in \mathcal{N}_k^i$, using the information received through local communication as in (11). These variables are then aggregated into the weighted average $\bar{x}_k^i \in \mathbb{R}^{\mathcal{P}}$, which acts as the reference point of the local proximal step used to compute the estimate $y_k^i \in \mathbb{R}^{\mathcal{P}}$ of the global solution, namely,

$$y_k^i = \operatorname{prox}_{f_k^i + g_k^i}^{1/\rho\eta_k^i}(\bar{x}_k^i), \quad \bar{x}_k^i = \frac{1}{\rho\eta_k^i} \sum_{j \in \mathcal{N}_k^i} x_k^{ij}. \quad (10)$$

However, the local proximal operator in (10) may not admit a closed form, since only g_k^i is guaranteed to be proximable. We therefore exploit the composite structure of $f_k^i + g_k^i$ to approximate (10) through a finite number $\tau \in \mathbb{N}_+$ of proximal gradient descent iterations with step-size $\pi \in \mathbb{R}_+$ as in eq. (12) [15, Sec. 28.5]. Each agent $i \in \mathcal{V}_k$ has $\mathcal{S}_k^i = \mathcal{N}_k^i$ and the time-varying set of labels $\mathcal{I}_k^i = \mathcal{S}_k^i \times \mathcal{P}$. The operator $\mathbb{T} : \mathbb{R}^{\mathcal{I}_k} \rightarrow \mathbb{R}^{\mathcal{I}_k}$ with $\mathcal{I}_k = \{(i, j, p) \mid i \in \mathcal{V}_k, j \in \mathcal{N}_k^i, p \in \mathcal{P}\}$ describing Algorithm 1 thus can be formalized block-wise as

$$x_k^{ij} = \mathbb{T}_k^{ij}(x_{k-1}) = \begin{cases} F_k^{ij}(x_{k-1}) + e_k^{ij} & \text{if } i \in \mathcal{R}_k \wedge j \in \mathcal{R}_k^i, \\ \rho y_k^{i,*} & \text{if } i \in \mathcal{A}_k \vee j \in \mathcal{A}_k^i, \end{cases}$$

¹As in Definitions 1.4, 1.21, and 8.1 of [15], respectively.

²The functions f_k^i are uniformly L -smooth if they are continuously differentiable and if their gradients ∇f_k^i are L -Lipschitz continuous for all $i \in \mathcal{V}_k$ and $k \in \mathbb{N}$. If the agents have different and time-varying local smoothness constants L_k^i , then L must satisfy $L \geq \sup_{k \in \mathbb{N}, i \in \mathcal{V}_k} L_k^i$. A function g_k^i is proximable if it has a closed-form/computationally cheap proximal, cf. [16].

Algorithm 1 Open ADMM (for composite problems)

Input: Relaxation parameter $\alpha \in (0, 1)$, penalty $\rho > 0$, number of local iterations $\tau \in \mathbb{N}_+$, step-size $\pi \in (0, 2/(L + \rho\eta))$
Output: Each agent returns y_k^i which is an approximate solution to the optimization problem in (8)

for $k = 0, 1, 2, \dots$ **each agent** $i \in \mathcal{V}_k$:

if $i \in \mathcal{A}_k$ // arriving agent

initializes the state to a local optimum

$$x_k^{ij} = \rho y_k^{i,*}, \quad y_{k-1}^i = y_k^{i,*}, \quad y_k^{i,*} \in \mathcal{Y}_k^{i,*}, \quad \forall j \in \mathcal{N}_k^i$$

else if $i \in \mathcal{R}_k$ // remaining agent

receives y_{k-1}^j, x_{k-1}^{ji} from each neighbor $j \in \mathcal{R}_k^i$
 updates the remaining state variable

$$x_k^{ij} = (1 - \alpha)x_{k-1}^{ij} - \alpha x_{k-1}^{ji} + 2\rho\alpha y_{k-1}^j, \quad \forall j \in \mathcal{R}_k^i \quad (11)$$

initializes the new states to a local optimum

$$x_k^{ij} = \rho y_k^{i,*} \text{ (see (9))}, \quad \forall j \in \mathcal{A}_k^i$$

end if

sets $r_0^i = y_{k-1}^i$ and approximates the proximal in (10) by:
for $t = 0, 1, \dots, \tau - 1$

$$r_{t+1}^i = \operatorname{prox}_{g_k^i}^{\pi} \left(r_t^i - \pi \left(\nabla f_k^i(r_t^i) + \rho\eta_k^i(r_t^i - \bar{x}_k^i) \right) \right) \quad (12)$$

end for

transmits $y_k^i = r_\tau^i, x_k^{ij}$ to each neighbor $j \in \mathcal{N}_k^i$
end for

where F_k^{ij} is the operator characterizing (10)–(11), and e_k^{ij} accounts for the inexact computation of the proximal (the discrepancy between (10) and (12)). We note that in practical scenarios, additional sources of errors can contribute to e_k^{ij} , e.g., inexact inter-agent communications due to quantization/compression [17]. These non-idealities are accommodated by assuming that the aggregate perturbation due to inexact local update and noisy communications is bounded.

Assumption 2. At all times $k \in \mathbb{N}$, each node $i \in \mathcal{V}_k$ updates its local variable y_k^i with an additive error $v_k^i \in \mathbb{R}^{\mathcal{P}}$ bounded by $\|v_k^i\| \leq \nu_u$ and receives information from each neighbor $j \in \mathcal{R}_k^i$ with an additive noise $v_k^{ij} \in \mathbb{R}^{\mathcal{P}}$ bounded by $\|v_k^{ij}\| \leq \nu_v$. Thus, the perturbation $e_k^{ij} = \alpha(v_k^{ij} + 2\rho v_{k-1}^j)$ on the state variable x_k^{ij} is bounded by $\|e_k^{ij}\| \leq \nu_v + 2\rho\nu_u =: \nu_e < \infty$.

Remark 1. By the definition of proximal in (2), computing (10) requires solving $\min_x f_k^i(x) + g_k^i(x) + \frac{\rho\eta_k^i}{2}\|x - \bar{x}_k^i\|^2$. The sum of the first and last term is $(\rho\eta_k^i)$ -strongly convex [15, Definition 10.7 and Proposition 10.8] and $(L + \rho\eta_k^i)$ -smooth. Consequently, the proximal gradient iteration (12) used to approximate (10) is contractive [18, Theorem 3.1] and, in turn, r_t^i converges linearly to the exact proximal, i.e., $\left\| y_k^i - \operatorname{prox}_{f_k^i + g_k^i}^{1/\rho\eta_k^i}(\bar{x}_k^i) \right\| \leq O(\zeta^\tau)$ with $\zeta = \max\{|1 - \pi(L + \rho\eta)|, |1 - \pi\rho \min_i \eta_k^i|\} \in (0, 1)$. Therefore, for any given bound $\nu_e \in (0, \infty)$, there always exists a sufficiently large τ so that Assumption 2 is satisfied.

Remark 2. In principle, r_0^i in (12) can be initialized arbitrarily. Inspired by [19], we set $r_0^i = y_{k-1}^i$, which creates a feedback loop that improves the approximation.

The main convergence result is given next. It shows, under mild assumptions on the graph topology given in Assumption 3, the global asymptotic open stability of the TSI of an OMAS executing Algorithm 1 and, as a consequence, the convergence of the agents' estimates to a consensus state on the set of optimal solutions.

Assumption 3. At all times $k \in \mathbb{N}_+$, the graph $\mathcal{G}_k$ of the OMAS is undirected and connected and the neighborhood size of every agent is uniformly bounded by $\bar{\eta}$, i.e., $\eta_k^i \leq \bar{\eta}$, $\forall i \in \mathcal{V}_k$.

Theorem 1. Consider an OMAS executing Algorithm 1 to distributedly solve the optimization problem (8) under Assumptions 1–3. If the OMAS is paracontractive with $\gamma \in (0, 1)$ and the departure process is bounded with $\beta \in (\gamma, 1)$, then the TSI is globally asymptotically open stable with radius

$$R = \frac{\rho(\sigma + \omega) + \nu_e}{1 - \gamma/\beta},$$

with linear rate $\theta = \gamma/\beta \in (0, 1)$. Also, the stacking $\tilde{y}_k = [\tilde{y}_k^1, \dots, \tilde{y}_k^{n_k}]^\top$ of the agents' estimates $\tilde{y}_k^i := \mathbb{1}_{\eta_k^i} \otimes y_k^i$ converges linearly to the consensus state on the optimal solutions $\tilde{\mathcal{C}}_k^* := \{\mathbb{1}_{|\mathcal{E}_k|} \otimes y_k^* \mid y_k^* \in \mathcal{Y}_k^*\}$ within radius $\Delta = (R + \nu_e)/\rho$.

Proof: The proof is given in Section III-B. ■

Remark 3. Exploiting [13, Proposition 1], one may ensure paracontractivity of the OMAS in practice by showing that the standard iteration F_k associated with the executed algorithm is both averaged (see [17, Definition 2]) and metrically subregular (see [17, Definition 1]).

For Open ADMM as in Algorithm 1, the iteration F_k is averaged because it is induced by the relaxed Peaceman–Rachford splitting applied to the dual of problem (8) [20], [21]. Metric subregularity, instead, is guaranteed when F_k admits affine lower and upper bounds [17, Proposition 1]. This condition is satisfied for several classes of optimization problems, including those with linear and quadratic costs whose proximal operators are affine, such as the problems we considered in Sections IV-A–IV-B (cf. [17, Sec. IV]).

Remark 4. The requirement $\beta > \gamma$ means that departures must be sufficiently slow compared with the contraction rate. As β approaches γ , the radius bound becomes conservative; this is a limitation of the analysis and does not necessarily reflect the actual convergence neighborhood observed in practice. Future work will investigate other metrics to overcome this limitation, e.g., infinity norms as in [14].

A. Stability of OMAS subject to additive error

Global asymptotic open stability is established for an OMAS whose state updates are affected by bounded additive perturbations. To provide sufficient conditions for stability in the sense of Definition 6, we impose limits on the variation of the TSI and on the processes by which agents join and leave. These limits are defined next.

Definition 7 (Bounded TSI). The TSI $\{\hat{\mathcal{X}}_k : k \in \mathbb{N}\}$ of an OMAS is said to have bounded variation if

$$\exists B \geq 0 : d_{\text{SH}}(\hat{\mathcal{X}}_k, \hat{\mathcal{X}}_{k-1}) \leq B\sqrt{|\mathcal{I}_k \cap \mathcal{I}_{k-1}|}, \quad \forall k \in \mathbb{N}.$$

Definition 8 (Bounded arrival process). The arrival process of an OMAS with TSI $\{\hat{\mathcal{X}}_k : k \in \mathbb{N}\}$ is said to be bounded if, letting $x_k^A \in \mathbb{R}^{\mathcal{I}_k \setminus \mathcal{I}_{k-1}}$ be the vector stacking the components of all arriving labels, it holds

$$\exists H \geq 0 : d(x_k^A, \hat{\mathcal{X}}_k) \leq H\sqrt{|\mathcal{I}_k \setminus \mathcal{I}_{k-1}|}, \quad \forall k \in \mathbb{N}.$$

Definition 9 (Bounded departure process). The departure process of an OMAS is said to be bounded if

$$\exists \beta \in (0, 1) : \sqrt{|\mathcal{I}_k|} \geq \beta\sqrt{|\mathcal{I}_{k-1}|}, \quad \forall k \in \mathbb{N}.$$

The main convergence result is given next.

Theorem 2. Given an OMAS, if:

- (a) it is paracontractive with $\gamma \in (0, 1)$;
 - (b) it admits a TSI with bounded variation $B \geq 0$;
 - (c) its arrival process is bounded with $H \geq 0$;
 - (d) its departure process is bounded with $\beta \in (\gamma, 1)$;
 - (e) the error is bounded by $\nu_e \geq 0$ as in Assumption 2;
- then the TSI is globally asymptotically open stable with radius

$$R = \frac{B + H + \nu_e}{1 - \gamma/\beta}. \quad (13)$$

Also, the convergence has linear rate $\theta = \gamma/\beta \in (0, 1)$ with

$$\frac{d(x_k, \hat{\mathcal{X}}_k)}{\sqrt{|\mathcal{I}_k|}} \leq \theta^k \frac{d(x_0, \hat{\mathcal{X}}_0)}{\sqrt{|\mathcal{I}_0|}} + (1 - \theta^k)R.$$

Proof: Let us split the state $x_k = [x_k^R; x_k^A]$ into two vectors such that $x_k^R \in \mathbb{R}^{\mathcal{I}_k \cap \mathcal{I}_{k-1}}$ is the vector of the remaining components and $x_k^A \in \mathbb{R}^{\mathcal{I}_k \setminus \mathcal{I}_{k-1}}$ is the vector of the new components. For any two consecutive steps it holds:

$$\begin{aligned} d(x_k, \hat{\mathcal{X}}_k) &= \sqrt{d^2(x_k^R, \hat{\mathcal{X}}_k) + d^2(x_k^A, \hat{\mathcal{X}}_k)} \leq \sqrt{d^2(x_k^R, \hat{\mathcal{X}}_k)} + \sqrt{d^2(x_k^A, \hat{\mathcal{X}}_k)} \\ &= d(x_k^R, \hat{\mathcal{X}}_k) + d(x_k^A, \hat{\mathcal{X}}_k) \stackrel{(i)}{\leq} d(x_k^R, \hat{\mathcal{X}}_k) + \sqrt{|\mathcal{I}_k|}H \\ &\stackrel{(ii)}{\leq} d(x_k^R, \hat{\mathcal{X}}_{k-1}) + d_{\text{SH}}(\hat{\mathcal{X}}_k, \hat{\mathcal{X}}_{k-1}) + \sqrt{|\mathcal{I}_k|}H \\ &\stackrel{(iii)}{\leq} d(x_k^R, \hat{\mathcal{X}}_{k-1}) + \sqrt{|\mathcal{I}_k|}(B + H) \\ &\stackrel{(iv)}{\leq} d(F_k(x_{k-1}) + e_k, \hat{\mathcal{X}}_{k-1}) + \sqrt{|\mathcal{I}_k|}(B + H) \\ &\stackrel{(v)}{\leq} d(F_k(x_{k-1}), \hat{\mathcal{X}}_{k-1}) + \|e_k\| + \sqrt{|\mathcal{I}_k|}(B + H) \\ &\stackrel{(vi)}{\leq} d(F_k(x_{k-1}), \hat{\mathcal{X}}_{k-1}) + \sqrt{|\mathcal{I}_k|}(B + H + \nu_e) \\ &\stackrel{(vii)}{\leq} \gamma d(x_{k-1}, \hat{\mathcal{X}}_{k-1}) + \sqrt{|\mathcal{I}_k|}(B + H + \nu_e), \end{aligned}$$

where (i) holds by assumption (c) and $|\mathcal{I}_k \setminus \mathcal{I}_{k-1}| \leq |\mathcal{I}_k|$; (ii) holds by triangle inequality in (1) for the shadow distance in Definition 2; (iii) holds by assumption (b) and $|\mathcal{I}_k \cap \mathcal{I}_{k-1}| \leq |\mathcal{I}_k|$; (iv) holds because the vector x_k^R is entirely contained in $F_k(x_{k-1}) + e_k$ by definition, where $e_k = [\dots, e_k^{ij}, \dots]^\top$ denotes the stacking of all error vectors for $i \in \mathcal{V}_k$, $j \in \mathcal{N}_k^i$; (v) holds by $\tilde{F}_k^{ij}(x_{k-1}) = F_k^{ij}(x_{k-1}) + e_k^{ij}$ and the triangle inequality; (vi) holds because the block components satisfy $\|e_k^{ij}\| \leq \nu_e$ by assumption (e) and the number of blocks is upper bounded by $|\mathcal{I}_k|$ as $e_k \in \mathbb{R}^{\mathcal{I}_k}$; (vii) holds by assumption (a). Using (d) and iterating over $k \in \mathbb{N}$ one gets

$$\frac{d(x_k, \hat{\mathcal{X}}_k)}{\sqrt{|\mathcal{I}_k|}} \leq \left(\frac{\gamma}{\beta}\right)^k \frac{d(x_0, \hat{\mathcal{X}}_0)}{\sqrt{|\mathcal{I}_0|}} + (B + H + \nu_e) \sum_{i=0}^{k-1} \left(\frac{\gamma}{\beta}\right)^i.$$

Since $\lim_{k \rightarrow \infty} (\gamma/\beta)^k = 0$ and $\lim_{k \rightarrow \infty} (1 - (\gamma/\beta)^k)/(1 - \gamma/\beta) = 1/(1 - \gamma/\beta)$, the proof is concluded by $\limsup_{k \rightarrow \infty} \frac{d(x_k, \hat{\mathcal{X}}_k)}{\sqrt{|\mathcal{I}_k|}} \leq \frac{B + H + \nu_e}{1 - \gamma/\beta} = R$. ■

B. Proof of Theorem 1

The first part of the proof consists in showing that conditions (a) – (e) of Theorem 2 hold: (a) Paracontractivity of the operator ruling the standard dynamics holds by assumption with $\gamma \in (0, 1)$; (b) Boundedness of the variation of the TSI holds by [13, Lemma 1] with $B = \rho\sigma$; (c) Boundedness of the arrival process holds by [13, Lemma 2] with $H = \rho\omega$; (d) Boundedness of the departure process holds by assumption with $\beta \in (\gamma, 1)$; (e) boundedness of the additive error holds by Assumption 2; then, substituting H and B into (13) gives the stated radius. The second part of the proof follows by:

$$\begin{aligned} d(\tilde{y}_k, \tilde{\mathcal{C}}_k^*) &= \sqrt{\sum_{i \in \mathcal{V}_k} \sum_{j \in \mathcal{N}_k^i} d^2(y_k^i, y_k^j)} = \sqrt{\sum_{i \in \mathcal{V}_k} \eta_k^i d^2(y_k^i, y_k^*)} \\ &\stackrel{(a)}{\leq} \sqrt{\sum_{i \in \mathcal{V}_k} \eta_k^i ((\rho \sqrt{\eta_k^i})^{-1} d(x_k^i, \hat{\mathcal{X}}_k) + \rho^{-1} \nu_e)^2} \\ &\stackrel{(b)}{\leq} \sqrt{\sum_{i \in \mathcal{V}_k} \rho^{-2} d^2(x_k^i, \hat{\mathcal{X}}_k) + \sum_{i \in \mathcal{V}_k} \rho^{-2} \eta_k^i \nu_e^2} \\ &\stackrel{(c)}{=} \rho^{-1} (d(x_k, \hat{\mathcal{X}}_k) + \sqrt{|\mathcal{E}_k|} \nu_e), \end{aligned}$$

where (b) holds by Minkowski's inequality, (c) holds by $\sum_{i \in \mathcal{V}_k} \eta_k^i = |\mathcal{E}_k|$, and (a) holds by

$$\begin{aligned} d(y_k^i, \mathcal{Y}_k^*) &= \inf_{y \in \mathcal{Y}_k^*} \|y_k^i - y\| \stackrel{(i)}{\leq} \|y_k^i - y_k^*\| \\ &\stackrel{(ii)}{\leq} \left\| \text{prox}_{f_k^i + g_k^i}^{1/\rho \eta_k^i} \left(\frac{1}{\rho \eta_k^i} A_{i,k}^\top x_k^i + u_k^i - \text{prox}_{f_k^i + g_k^i}^{1/\rho \eta_k^i} \left(\frac{1}{\rho \eta_k^i} A_{i,k}^\top \hat{x}_k^i \right) \right) \right\| \\ &\stackrel{(iii)}{\leq} \left\| \frac{1}{\rho \eta_k^i} A_{i,k}^\top (x_k^i - \hat{x}_k^i) \right\| + \rho^{-1} \nu_e \stackrel{(iv)}{\leq} \left\| \frac{1}{\rho \eta_k^i} A_{i,k}^\top \right\| \|x_k^i - \hat{x}_k^i\| + \rho^{-1} \nu_e \\ &\stackrel{(v)}{=} (\rho \sqrt{\eta_k^i})^{-1} d(x_k^i, \hat{\mathcal{X}}_k) + \rho^{-1} \nu_e, \end{aligned}$$

where (i) holds since $y_k^* \in \mathcal{Y}_k^*$; (ii) holds by eqs. (10), where $A_{i,k}^\top$ is the matrix computing the sums of neighbors' states and $u_k^i \in \mathbb{R}^p$ is the error on the local variable as defined in Assumption 2; (iii) follows by the non-expansiveness of the proximal operator and by Assumption 2 yielding $\|u_k^i\| \leq \nu_u \leq (2\rho)^{-1} \nu_e \leq \rho^{-1} \nu_e$; (iv) holds by choosing $\hat{x}_k = \arg\inf_{y \in \hat{\mathcal{X}}_k} \|x_k - y\|$; (v) holds because matrix $\frac{1}{\eta_k^i} A_{i,k}^\top$ is column-stochastic and its rows sum up to $1/\eta_k^i$. This means that the linear convergence of $\tilde{y}_k$ to a neighborhood of $\tilde{\mathcal{C}}_k^*$ is implied by that of x_k to a neighborhood of $\hat{\mathcal{X}}_k$, which follows from the first part of the proof, up to the convergence radius:

$$\limsup_{k \rightarrow \infty} \frac{d(\tilde{y}_k, \tilde{\mathcal{C}}_k^*)}{\sqrt{|\mathcal{P}||\mathcal{E}_k|}} \leq \limsup_{k \rightarrow \infty} \frac{d(x_k, \hat{\mathcal{X}}_k) + \sqrt{|\mathcal{E}_k|} \nu_e}{\rho \sqrt{|\mathcal{P}||\mathcal{E}_k|}} \leq \frac{R + \nu_e}{\rho}.$$

IV. APPLICATION TO DISTRIBUTED LEARNING WITH LOCAL CONSTRAINTS

This section showcases the performance of Open ADMM applied to composite (un)supervised learning problems, which, in real scenarios, must be solved in open networks because the participating data holders (e.g., devices, sensors, or local processors) may join or leave the network over time due to intermittent availability, mobility, failures, or communication constraints. We focus on problems with local constraints³:

$$\min_{y_i \in \mathcal{Y}_k^i} \sum_{i \in \mathcal{V}_k} f_k^i(y_i) \quad \text{s.t.} \quad y_i = y_j \quad \forall (i, j) \in \mathcal{E}_k, \quad (14)$$

where $\mathcal{Y}_k^i \subset \mathbb{R}^p$, with $i \in \mathcal{V}_k$, are nonempty, closed, convex sets such that $\bigcap_{i \in \mathcal{V}_k} \mathcal{Y}_k^i \neq \emptyset$ for all $k \in \mathbb{N}$. This problem is a special case of (8) with the indicator function $g_k^i = \iota_{\mathcal{Y}_k^i}$, such that $g_k^i(y) = 0$ if $y \in \mathcal{Y}_k^i$ and $g_k^i(y) = +\infty$ otherwise. The performance will be evaluated in terms of:

- (i) optimality error $\text{err}_k := d(\bar{y}_k, \mathcal{Y}_k^*)$, $\bar{y}_k = \sum_{i \in \mathcal{V}_k} y_k^i / |\mathcal{V}_k|$;
- (ii) constraint violation $\text{cv}_k := d(\bar{y}_k, \bigcap_{i \in \mathcal{V}_k} \mathcal{Y}_k^i)$.

A. Supervised learning: linear regression

We apply Open ADMM to solve distributed linear regression problems over open networks of agents given next:

$$f_k^i(y) = \frac{1}{2} \|A_k^i{}^\top y - b_k^i\|^2, \quad \mathcal{Y}_k^i = [\underline{u}_k^i, \bar{u}_k^i], \quad \forall i \in \mathcal{V}_k,$$

where $A_k^i = [\dots, a_{k,h}^i, \dots] \in \mathbb{R}^{p \times m_k^i}$ is the matrix collecting all the feature vectors, $b_k^i = [\dots, b_{k,h}^i, \dots]^\top \in \mathbb{R}^{m_k^i}$ is the vector collecting the labels, and $\underline{u}_k^i \in [-20, -15]$, $\bar{u}_k^i \in [-10, -5]$ define the box constraints. The data are randomly generated using `scikit-learn`. This problem belongs to the class of *supervised learning*, where each agent's local dataset has time-varying cardinality $m_k^i \in \mathbb{N}_+$ and contains pairs $(a_{k,h}^i, b_{k,h}^i)$ where $a_{k,h}^i \in \mathbb{R}^p$ are $p \in \mathbb{N}_+$ features and $b_{k,h}^i \in \mathbb{R}$ is the associated label. The goal is to learn a model y that predicts the label $b_{k,h}^i$ from the sample $a_{k,h}^i$. The number of arriving and departing agents is drawn from independent Poisson distributions with the same mean.

We start by displaying in Table I the performance of Open ADMM with different values of the local number of iterations τ , in terms of asymptotic optimality error and constraint violation. Interestingly, as τ increases err and cv decrease, plateauing as $\tau \geq 5$. This is due to the fact that the main source of suboptimality when $\tau \geq 5$ are the variations in the network, rather than the inexact proximal computations.

While the results achieved by Open ADMM are already satisfying, the particular structure of the constraints inspired the design of a variant algorithm with better performance. Since the local constraint sets are box constraints $\mathcal{Y}_k := \bigcap_{i \in \mathcal{V}_k} \mathcal{Y}_k^i = [\max_{i \in \mathcal{V}_k} \underline{u}_k^i, \min_{i \in \mathcal{V}_k} \bar{u}_k^i] =: [\underline{u}_k, \bar{u}_k]$, they can be estimated through max-consensus protocols, such as OSTDMC [14] or Open ADMM itself [13, Sec. IV.A]. An improvement can thus be obtained by preceding with a max-consensus step on $(\underline{u}_k^i, -\bar{u}_k^i)$ to obtain a tighter estimate $\hat{\mathcal{Y}}_k^i = [\hat{\underline{u}}_k^i, \hat{\bar{u}}_k^i]$ of $\mathcal{Y}_k$; in our simulations, this max-consensus step is implemented using Open ADMM as in [13, Sec. IV.A]. Therefore, during the local step (12) of Open ADMM, the agents apply projections onto $\hat{\mathcal{Y}}_k^i$ rather than $\mathcal{Y}_k^i$. Table I shows that this modified algorithm manages to outperform Open ADMM by one order of magnitude, both in terms of optimality error and constraint violation.

TABLE I
ASYMPTOTIC VALUES OF OPTIMALITY ERROR AND CONSTRAINT VIOLATION FOR DIFFERENT VALUES OF τ .

Open ADMM	τ	1	2	5	10	1000
standard	err	0.2371	0.2368	0.2367	0.2367	0.2367
	cv	0.2594	0.2585	0.2587	0.2587	0.2587
with max-consensus	err	0.0126	0.0125	0.0125	0.0125	0.0125
	cv	0.0498	0.0499	0.0499	0.0499	0.0499

³See code <https://github.com/nicola-bastianello/open-admm>

B. Unsupervised learning: distributed PCA

We apply Open ADMM to solve a distributed principal component analysis (PCA) problem over an open network of agents given next:

$$f_k^i(Y) = -\text{tr}(YC_{i,k}), \quad \mathcal{Y}_k^i = \mathcal{F}_{p,q}, \quad (15)$$

where $Y \in \mathbb{R}^{p \times p}$ is the optimization variable, $C_{i,k} = \frac{1}{m_{i,k}} A_k^i A_k^{i\top} \in \mathbb{R}^{p \times p}$ is the local correlation matrix, $A_k^i = [\cdots, a_{k,h}^i, \cdots] \in \mathbb{R}^{p \times m_k^i}$ is the matrix collecting all the feature vectors, $\mathcal{F}_{p,q} := \{M \in \mathbb{R}^{p \times p} \mid 0 \preceq M \preceq 1, \text{tr}(M) = q\}$ is the *Fantope* set [22, Sec. 2.3.2], with $M \preceq a$ denoting a symmetric matrix with eigenvalues less or equal than $a \in \mathbb{R}$, and $\text{tr}(M) = \sum_{\ell=1}^p M_{\ell\ell}$. This problem belongs to the class of *unsupervised learning*, where each agent holds a time-varying dataset of unlabeled feature vectors $a_{k,h}^i \in \mathbb{R}^p$, and the goal is to learn a projection matrix $U_{q,k} \in \mathbb{R}^{p \times q}$, with $q < p$, maximizing the variance of the projected data [23]; each agent constructs $U_{q,k}$ from the q leading eigenvectors of its local solution Y_i . Datasets of heterogeneous size $m_i \in [50, 200]$ are randomly generated with a latent $2D$ subspace of five clusters embedded in a $p = 50$ feature space (SNR= 25), and PCA is performed with $q = 2$.

Figure 1 reports, on a logarithmic scale, the normalized tracking error err_k together with the number of agents n_k . To assess the impact of agent arrivals and departures, simulations are run with join/leave events drawn from Poisson distributions with equal mean $\lambda^{\text{join}} = \lambda^{\text{leave}} = \lambda$, considering $\lambda \in \{0.1, 1, 100\}$. As expected, larger λ yields larger err_k due to stronger fluctuations in network size; nevertheless, after the initial transient, err_k settles within a bounded distance from the optimal solution, in agreement with the theoretical guarantees. Finally, Figure 2 shows a two-dimensional projection of the data for five randomly selected agents at a representative time under the worst-case setting $\lambda = 100$, confirming that the learned projection preserves the five-cluster structure.

V. CONCLUSIONS

This paper studied distributed *composite optimization* for OMAS by introducing a practical and effective variant of Open ADMM [13] in which local composite proximal steps




Fig. 1. Open ADMM applied to an unsupervised learning problem in an open network with $\lambda^{\text{join}} = \lambda^{\text{leave}} = \lambda \in \{0.1, 1, 100\}$.




Fig. 2. Two-dimensional latent-space projection of the datasets for five randomly selected agents in the network at a random time during the simulation, under the worst-case configuration $\lambda = 100$.

are replaced by a finite number of proximal-gradient iterations. A reusable analysis tool for algorithm design in open networks is developed by establishing a general convergence result for paracontractive OMAS under bounded additive perturbations, which yields a linear tracking rate and a steady-state radius. Future work will focus on asynchronous communications, nonconvex learning settings, and robustness to attacks.

REFERENCES

- [1] N. K. Dhingra, S. Z. Khong, and M. R. Jovanovic, "The proximal augmented lagrangian method for nonsmooth composite optimization," *IEEE Transactions on Automatic Control*, vol. 64, no. 7, 2019.
- [2] D. Yuan, Y. Hong, D. W. C. Ho, and S. Xu, "Distributed mirror descent for online composite optimization," *IEEE Transactions on Automatic Control*, vol. 66, no. 2, 2021.
- [3] R. Dixit, A. S. Bedi, and K. Rajawat, "Online learning over dynamic graphs via distributed proximal gradient algorithm," *IEEE Transactions on Automatic Control*, vol. 66, no. 11, 2021.
- [4] N. Bastianello and E. Dall'Anese, "Distributed and inexact proximal gradient method for online convex optimization," in *2021 European Control Conference (ECC)*, Delft, Netherlands, 2021.
- [5] X. Li, L. Xie, and N. Li, "A survey on distributed online optimization and online games," *Annual Reviews in Control*, vol. 56, 1, 2023.
- [6] R. Hou, Y. Yu, and X. Li, "Dynamic regret for distributed online composite optimization," *IEEE Transactions on Automatic Control*, vol. 70, no. 5, 2024.
- [7] Z. Yu, D. W. Ho, D. Yuan, and Z. Shi, "Distributed composite optimization with distributed composite block mirror descent," *IEEE Transactions on Automatic Control*, 2026.
- [8] H. Tian, X. Li, Z. Wu, and S. Magnússon, "Communication-efficient online federated composite optimization," *Automatica*, vol. 183, 2026.
- [9] Y. Zhou, L. Guo, X. Shi, and J. Cao, "A distributed proximal alternating direction multiplier method for multiblock nonsmooth composite optimization," *IEEE Transactions on Control of Network Systems*, vol. 12, no. 1, 2025.
- [10] M. Franceschelli and P. Frasca, "Stability of open multiagent systems and applications to dynamic consensus," *IEEE Transactions on Automatic Control*, vol. 66, no. 5, 2020.
- [11] D. Deplano, M. Franceschelli, and A. Giua, "Stability of paracontractive open multi-agent systems," in *63rd IEEE Conference on Decision and Control*, 2024.
- [12] R. Vizuete, C. Monnayer de Galland, P. Frasca, E. Panteley, and J. M. Hendrickx, "Trends and questions in open multi-agent systems," in *Hybrid and Networked Dynamical Systems: Modeling, Analysis and Control*, Springer, 2024.
- [13] D. Deplano, N. Bastianello, M. Franceschelli, and K. H. Johansson, "Optimization and learning in open multi-agent systems," *IEEE Transactions on Automatic Control*, 2026, to appear.
- [14] D. Deplano, M. Franceschelli, and A. Giua, "Distributed tracking of network size, diameter, radius, and node eccentricities in open multi-agent systems," *IEEE Transactions on Automatic Control*, 2025.
- [15] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, 2nd ed. Cham, Switzerland: Springer, 2017.
- [16] N. Parikh and S. Boyd, "Proximal algorithms," *Foundations and Trends® in Optimization*, vol. 1, no. 3, 2014.
- [17] N. Bastianello, D. Deplano, M. Franceschelli, and K. H. Johansson, "Robust online learning over networks," *IEEE Transactions on Automatic Control*, vol. 70, no. 2, 2025.
- [18] A. B. Taylor, J. M. Hendrickx, and F. Glineur, "Exact worst-case convergence rates of the proximal gradient method for composite convex minimization," *Journal of Optimization Theory and Applications*, vol. 178, no. 2, 2022.
- [19] X. Ren, N. Bastianello, K. H. Johansson, and T. Parisini, "Communication-efficient stochastic distributed learning," *IEEE Transactions on Automatic Control*, 2026.
- [20] N. Bastianello, R. Carli, L. Schenato, and M. Todescato, "Asynchronous distributed optimization over lossy networks via relaxed admm: Stability and linear convergence," *IEEE Transactions on Automatic Control*, vol. 66, no. 6, 2020.
- [21] P. Giselsson and S. Boyd, "Linear convergence and metric selection for douglas-rachford splitting and admm," *IEEE Transactions on Automatic Control*, vol. 62, no. 2, 2016.
- [22] J. Dattorro, *Convex optimization & Euclidean distance geometry*. Lulu.com, 2010.

- [23] X. Chang, F. Nie, Y. Yang, C. Zhang, and H. Huang, "Convex sparse pca for unsupervised feature learning," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 11, no. 1, 2016.