A Riemannian gossip approach to decentralized matrix completion

Bamdev Mishra
Amazon Development Centre India,
Bangalore, India
bamdevm@amazon.com

Hiroyuki Kasai
The University of Electro-Communications
Tokyo, Japan
kasai@is.uec.ac.jp

Atul Saroop
Amazon Development Centre India,
Bangalore, India
asaroop@amazon.com

Abstract

In this paper, we propose novel gossip algorithms for the low-rank decentralized matrix completion problem. The proposed approach is on the Riemannian Grassmann manifold that allows *local* matrix completion by different agents while achieving asymptotic consensus on the *global* low-rank factors. The resulting approach is scalable and parallelizable. Our numerical experiments show the good performance of the proposed algorithms on various benchmarks.

1 Introduction

The problem of low-rank matrix completion amounts to completing a matrix from a small number of entries by assuming a low-rank model for the matrix. The problem has many applications in control systems and system identification [1], collaborative filtering [2], and information theory [3], to name a just few. Consequently, it has been a topic of great interest and there exist many large-scale implementations for both *batch* [4, 5, 6, 7, 8, 9, 10] and *online* scenarios that focus on parallel and stochastic implementations [11, 12, 13, 14].

In this paper, we are interested in a *decentralized* setting, where we divide the matrix completion problem into smaller subproblems that are solved by many agents locally while simultaneously enabling them to arrive at a *consensus* that solves the full problem [15]. The recent paper [15] proposes a particular decentralized framework for matrix completion by exploiting the algorithm proposed in [6]. It, however, requires an inexact dynamic consensus step at every iteration. We relax this by proposing a novel formulation that combines together a weighted sum of completion and consensus terms. Additionally, in order to minimize the communication overhead between the agents, we constrain each agent to communicate with *only* one other agent as in the *gossip* framework [16]. One motivation is that this addresses privacy concerns of sharing sensitive data [15]. Another motivation is that the gossip framework is robust to scenarios where certain agents may be inactive at certain time slots, e.g., consider each agent to be a computing machine. We propose a preconditioned variant that is particularly well suited for ill-conditioned instances. Additionally, we also propose a parallel variant that allows to exploit parallel computational architectures. All the variants come with asymptotic convergence guarantees. To the best of our knowledge, this is the first work that exploits the gossip architecture for solving the decentralized matrix completion problem.

1

The organization of the paper is as follows. In Section 2, we discuss the decentralized problem setup and propose a novel problem formulation. In Section 3, we discuss the proposed stochastic gradient gossip algorithm for the matrix completion problem. A preconditioned variant of the Riemannian gossip algorithm is motivated in Section 3.3. Additionally, we discuss a way to parallelize the proposed algorithms in Section 3.4. Numerical comparisons in Section 4 show that the proposed algorithms compete effectively with state-of-the-art on various benchmarks. The Matlab codes for the proposed algorithms are available at `https://bamdevmishra.com/codes/gossipMC/`.

2 Decentralized matrix completion

The matrix completion problem is formulated as

$$\min_{\mathbf{X} \in \mathbb{R}^{m \times n}} \quad \frac{1}{2} \| \mathcal{P}_\Omega(\mathbf{X}) - \mathcal{P}_\Omega(\mathbf{X}^\star) \|_F^2 \tag{1}$$
$$\text{subject to} \quad \text{rank}(\mathbf{X}) = r,$$

where $\mathbf{X}^\star \in \mathbb{R}^{n \times m}$ is a matrix whose entries are known for indices if they belong to the subset $(i,j) \in \Omega$ and Ω is a subset of the complete set of indices $\{(i,j) : i \in \{1,...,m\} \text{ and } j \in \{1,...,n\}\}$. The operator $\mathcal{P}_\Omega(\mathbf{X}_{ij}) = \mathbf{X}_{ij}$ if $(i,j) \in \Omega$ and $\mathcal{P}_\Omega(\mathbf{X}_{ij}) = 0$ otherwise is called the orthogonal sampling operator and is a mathematically convenient way to represent the subset of known entries. The rank constraint parameter r is usually set to a low value, i.e., $\ll (m,n)$ that implies that we seek low-rank completion. A way to handle the rank constraint in (1) is by a fixed-rank matrix parameterization. In particular, we use $\mathbf{X} = \mathbf{U}\mathbf{W}^T$, where $\mathbf{U} \in \text{St}(r,m)$ and $\mathbf{W} \in \mathbb{R}^{n \times r}$, where $\text{St}(r,m)$ is the set of $m \times r$ *orthonormal* matrices, i.e., the columns are orthonormal. The interpretation is that $\mathbf{U}$ captures the dominant column space of $\mathbf{X}$ and $\mathbf{W}$ captures the *weights* [17]. Consequently, the optimization problem (1) reads

$$\min_{\mathbf{U} \in \text{St}(r,m)} \min_{\mathbf{W} \in \mathbb{R}^{n \times r}} \| \mathcal{P}_\Omega(\mathbf{U}\mathbf{W}^T) - \mathcal{P}_\Omega(\mathbf{X}^\star) \|_F^2. \tag{2}$$

The *inner* least-squares optimization problem in (2) is solved in closed form by exploiting the least-squares structure of the cost function to obtain the optimization problem

$$\min_{\mathbf{U} \in \text{St}(r,m)} \quad \frac{1}{2} \| \mathcal{P}_\Omega(\mathbf{U}\mathbf{W_U}^T) - \mathcal{P}_\Omega(\mathbf{X}^\star) \|_F^2 \tag{3}$$

in $\mathbf{U}$, where $\mathbf{W_U}$ is the solution to the inner optimization problem $\min_{\mathbf{W} \in \mathbb{R}^{n \times r}} \| \mathcal{P}_\Omega(\mathbf{U}\mathbf{W}^T) - \mathcal{P}_\Omega(\mathbf{X}^\star) \|_F^2$. (The cost function in (3) may be discontinuous at points $\mathbf{U}$ where $\mathbf{W_U}$ is non-unique [18]. This is handled effectively by adding a regularization term $\| \mathbf{X} \|_F^2$ to (1).)

The problem (3) requires handling the entire incomplete matrix $\mathbf{X}^\star$ at all steps of optimization. This is memory intensive and computationally heavy, especially in large-scale instances. To relax this constraint, we distribute the task of solving the problem (3) among N agents, which perform certain computations independently. To this end, we partition the incomplete matrix $\mathbf{X}^\star = [\mathbf{X}_1^\star, \mathbf{X}_2^\star, \ldots, \mathbf{X}_N^\star]$ along the columns such that the size of $\mathbf{X}_i^\star$ is $m \times n_i$ with $\sum n_i = n$ for $i = \{1, 2, \ldots, N\}$. Each agent i has knowledge of the incomplete matrix $\mathbf{X}_i^\star$ and its *local* set of indices Ω_i of known entries. We also partition the weight matrix $\mathbf{W}$ as $\mathbf{W}^T = [\mathbf{W}_1^T, \mathbf{W}_2^T, \ldots, \mathbf{W}_N^T]$ such that the matrix $\mathbf{W}_i$ has size $n_i \times r$. A straightforward reformulation of (3) is

$$\sum_i \min_{\mathbf{U} \in \text{St}(r,m), \mathbf{W}_i \in \mathbb{R}^{n_i \times r}} \quad \frac{1}{2} \| \mathcal{P}_{\Omega_i}(\mathbf{U}\mathbf{W}_i^T) - \mathcal{P}_{\Omega_i}(\mathbf{X}_i^\star) \|_F^2$$
$$= \min_{\mathbf{U} \in \text{St}(r,m)} \quad \frac{1}{2} \sum_i \underbrace{\| \mathcal{P}_{\Omega_i}(\mathbf{U}\mathbf{W}_{i\mathbf{U}}^T) - \mathcal{P}_{\Omega_i}(\mathbf{X}_i^\star) \|_F^2}_{\text{problem handled by agent } i}, \tag{4}$$

where $\mathbf{W}_{i\mathbf{U}}$ is the least-squares solution to $\min_{\mathbf{W}_i \in \mathbb{R}^{n_i \times r}} \| \mathcal{P}_{\Omega_i}(\mathbf{U}\mathbf{W}_i^T) - \mathcal{P}_{\Omega_i}(\mathbf{X}_i^\star) \|_F^2$, which can be computed by agent i independently of other agents.

Although the computational workload gets distributed among the agents in the problem formulation (4), all agents require the knowledge of $\mathbf{U}$ (to compute matrices $\mathbf{W}_{i\mathbf{U}}$). To circumvent this issue, instead of one shared matrix $\mathbf{U}$ for all agents, each agent i stores a local copy $\mathbf{U}_i$, which it then updates based on information from its *neighbors*. For minimizing the communication overhead

between agents, we additionally put the constraint that at any time slot only two agents communicate, i.e, each agent has exactly only one neighbor. This is the basis of the gossip framework [16]. In standard gossip framework, at a time slot, an agent is randomly assigned one neighbor [16]. However, to motivate the various ideas in this paper and to keep the exposition simple, we fix the agents network topology, i.e., each agent is preassigned a unique neighbor. (In Section 3.5, we show how to deal with random assignments of neighbors.) To this end, the agents are numbered according to their proximity, e.g., for $i \leqslant N - 1$, agents i and $i + 1$ are neighbors. Equivalently, agents 1 and 2 are neighbors and can communicate. Similarly, agents 2 and 3 communicate, and so on. This communication between the agents allows to reach a *consensus* on $\mathbf{U}_i$. Specifically, it suffices that the *column spaces* of all $\mathbf{U}_i$ converge. (The precise motivation and formulation are in Section 3.) Our proposed decentralized matrix completion problem formulation is

$$
\min_{\mathbf{U}_1,\ldots,\mathbf{U}_N \in \mathrm{St}(r,m)} \quad \frac{1}{2} \sum_i \underbrace{\|\mathcal{P}_{\Omega_i}(\mathbf{U}_i \mathbf{W}_{i\mathbf{U}_i}^T) - \mathcal{P}_{\Omega_i}(\mathbf{X}_i^\star)\|_F^2}_{\text{completion task handled by agent } i}
$$
$$
+ \frac{\rho}{2} \underbrace{(d_1(\mathbf{U}_1, \mathbf{U}_2)^2 + d_2(\mathbf{U}_2, \mathbf{U}_3)^2 + \ldots + d_{N-1}(\mathbf{U}_{N-1}, \mathbf{U}_N)^2)}_{\text{consensus}}, \tag{5}
$$

where d_i is a certain distance measure (defined in Section 3) between $\mathbf{U}_i$ and $\mathbf{U}_{i+1}$ for $i \leqslant N - 1$, minimizing which forces $\mathbf{U}_i$ and $\mathbf{U}_{i+1}$ to an "average" point (specifically, average of the column spaces). $\rho \geqslant 0$ is a parameter that trades off matrix completion with consensus. Here $\mathbf{W}_{i\mathbf{U}_i}$ is the solution to the optimization problem $\min_{\mathbf{W}_i \in \mathbb{R}^{n_i \times r}} \|\mathcal{P}_{\Omega_i}(\mathbf{U}_i \mathbf{W}_i^T) - \mathcal{P}_{\Omega_i}(\mathbf{X}_i^\star)\|_F^2$.

In standard gossip framework, the aim is to make the agents converge to a common point, e.g, minimizing only the consensus term in (5). In our case, we additionally need the agents to perform certain tasks, e.g., minimizing the completion term in (5), which motivates the *weighted* formulation (5). For a large value of ρ, the consensus term in (5) dominates, minimizing which allows the agents to arrive at consensus. For $\rho = 0$, the optimization problem (5) solves N independent completion problems and there is no consensus. For a sufficiently large value of ρ, the problem (5) achieves the goal of matrix completion along with consensus.

3 The Riemannian gossip algorithm

It should be noted that the optimization problem (3), and similarly (4), only depends on the *column space* of $\mathbf{U}$ rather than $\mathbf{U}$ itself [9, 11]. Equivalently, the cost function in (3) remains constant under the transformation $\mathbf{U} \mapsto \mathbf{U}\mathbf{O}$ for all orthogonal matrices $\mathbf{O}$ of size $r \times r$. Mathematically, the column space of $\mathbf{U}$ is captured by the set, called the *equivalence class*, of matrices

$$
[\mathbf{U}] := \{\mathbf{U}\mathbf{O} : \mathbf{O} \text{ is a } r \times r \text{ orthogonal matrix}\}. \tag{6}
$$

The set of the equivalence classes is called the *Grassmann* manifold, denoted by $\mathrm{Gr}(r, m)$, which is the set of r-dimensional subspaces in $\mathbb{R}^m$ [19]. The Grassmann manifold $\mathrm{Gr}(r, m)$ is identified with the quotient manifold $\mathrm{St}(r, m)/\mathcal{O}(r)$, where $\mathcal{O}(r)$ is the orthogonal group of $r \times r$ matrices [19].

Subsequently, the problem (3), and similarly (4), is on the Grassmann manifold $\mathrm{Gr}(r, m)$ and not on $\mathrm{St}(r, m)$. However, as $\mathrm{Gr}(r, m)$ is an abstract quotient space, numerical optimization algorithms are implemented with matrices $\mathbf{U}$ on $\mathrm{St}(r, m)$, but conceptually, optimization is on $\mathrm{Gr}(r, m)$. It should be stated that the Grassmann manifold has the structure of a *Riemannian* manifold and optimization on the Grassmann manifold is a well studied topic in literature. Notions such as the Riemannian gradient (first order derivatives of a cost function), geodesic (shortest distance between elements), and logarithm mapping (capturing "difference" between elements) have closed-form expressions [19].

If x is an element of a Riemannian compact manifold $\mathcal{M}$, then the decentralized formulation (5) boils down to the form

$$
\min_{x_1,\ldots,x_N \in \mathcal{M}} \quad \sum_i f_i(x_i)
$$
$$
+ \frac{\rho}{2} \underbrace{(d_1(x_1, x_2)^2 + d_2(x_2, x_3)^2 + \ldots + d_{N-1}(x_{N-1}, x_N)^2)}_{\text{consensus}}, \tag{7}
$$

Table 1: Proposed online gossip algorithm for (7)

1. At each time slot t, pick an agent $i \leqslant N-1$ randomly with uniform probability.
2. Compute the Riemannian gradients $\mathrm{grad}_{x_i} f_i$, $\mathrm{grad}_{x_{i+1}} f_{i+1}$, $\mathrm{grad}_{x_i} d_i$, and $\mathrm{grad}_{x_{i+1}} d_i$ with the matrix representations

$$
\begin{aligned}
\mathrm{Grad}_{x_i} f_i &= (\mathcal{P}_{\Omega_i}(\mathbf{U}_i \mathbf{W}_{i\mathbf{U}_i}^T) - \mathcal{P}_{\Omega_i}(\mathbf{X}_i^\star)) \mathbf{W}_{i\mathbf{U}_i} \\
\mathrm{grad}_{x_i} f_i &= \mathrm{Grad}_{x_i} f_i - \mathbf{U}_i(\mathbf{U}_i^T \mathrm{Grad}_{x_i} f_i) \\
\mathrm{grad}_{x_i} d_i &= -\mathrm{Log}_{x_i}(x_{i+1}) \\
\mathrm{grad}_{x_{i+1}} d_i &= -\mathrm{Log}_{x_{i+1}}(x_i),
\end{aligned}
$$

   where $\mathbf{U}_i$ is the matrix representation of x_i. $\mathrm{Log}_{x_i}(x_{i+1})$ is the *logarithm* mapping, which is defined as

$$
\mathrm{Log}_{x_i}(x_{i+1}) = \mathbf{P}\arctan(\mathbf{S})\mathbf{Q}^T,
$$

   where $\mathbf{PSQ}^T$ is the rank-r singular value decomposition of $(\mathbf{U}_{i+1} - \mathbf{U}_i(\mathbf{U}_i^T \mathbf{U}_{i+1}))$ $(\mathbf{U}_i^T \mathbf{U}_{i+1})^{-1}$. The $\arctan(\cdot)$ operation is only on the diagonal entries. It should be noted that the Riemannian gradient of the Riemannian distance is the negative logarithm mapping [21].
3. Given a stepsize γ_t, update x_i and x_{i+1} as

$$
\begin{aligned}
x_{i+} &= \mathrm{Exp}_{x_i}(-\gamma_t(\alpha_i \mathrm{grad}_{x_i} f_i + \rho \mathrm{grad}_{x_i} d_i)) \\
x_{i+1+} &= \mathrm{Exp}_{x_{i+1}}(-\gamma_t(\alpha_{i+1} \mathrm{grad}_{x_{i+1}} f_{i+1} + \rho \mathrm{grad}_{x_{i+1}} d_i)),
\end{aligned}
$$

   where $\mathbf{U}_i$ is the matrix representation of x_i and $\alpha_i = 1$ if $i = \{1, N\}$, else $\alpha_i = 0.5$. $\mathrm{Exp}_{x_i}(\xi_{x_i}) = \mathbf{U}_i \mathbf{V}\cos(\Sigma)\mathbf{V}^T + \mathbf{W}\sin(\Sigma)\mathbf{V}^T$ is the *exponential* mapping and $\mathbf{W}\Sigma\mathbf{V}^T$ is the rank-r singular value decomposition of ξ_{x_i}. The $\cos(\cdot)$ and $\sin(\cdot)$ operations are only on the diagonal entries.

where $x_i = [\mathbf{U}_i]$ with matrix representation $\mathbf{U}_i \in \mathrm{St}(r, m)$, $\mathcal{M} = \mathrm{Gr}(r, m) = \mathrm{St}(r, m)/\mathcal{O}(r)$, $f_i : \mathcal{M} \to \mathbb{R}$ is a continuous function, and $d_i : \mathcal{M} \times \mathcal{M} \to \mathbb{R}$ is the Riemannian geodesic distance between x_i and x_{i+1}. Here $[\mathbf{U}_i]$ is the equivalence class defined in (6). The Riemannian distance d_i captures the distance between the subspaces $[\mathbf{U}_i]$ and $[\mathbf{U}_{i+1}]$. Minimizing only the consensus term in (7) is equivalent to computing the *Karcher mean* of N subspaces [20, 21].

We exploit the stochastic gradient descent setting framework proposed by Bonnabel [20] for solving (7), which is an optimization problem on the Grassmann manifold. In particular, we exploit the stochastic gradient algorithm in the gossip framework. To keep the analysis simple, we predefine the topology on the agents network. Following [20, Section 4.4], we make the following assumptions.

A1 Agents i and $i+1$ are neighbors for all $i \leqslant N-1$.

A2 At each time slot, say t, we pick an agent $i \leqslant N-1$ randomly with uniform probability. This means that we also pick agent $i+1$ (the neighbor of agent i). Subsequently, agents i and $i+1$ update x_i and x_{i+1}, respectively, by taking a *gradient descent step* with stepsize γ_t on $\mathcal{M}$. The stepsize sequence satisfies the standard conditions, i.e., $\sum \gamma_t^2 < \infty$ and $\sum \gamma_t = +\infty$ [20, Section 3].

Each time we pick an agent $i \leqslant N-1$, we equivalently also pick its neighbor $i+1$. Subsequently, we need to update both of them by taking a gradient descent step based on $f_i(x_i) + f_{i+1}(x_{i+1}) + \rho d_i(x_i, x_{i+1})^2/2$. Repeatedly updating the agents in this fashion is a stochastic process.

It should be noted that because of the particular topology and sampling that we assume (in **A1** and **A2**), on an average x_2 to x_{N-1} are updated *twice* the number of times x_1 and x_N are updated. For example, if $N = 3$, then **A1** and **A2** lead to solving (in expectation) the problem $\min_{x_1, x_2, x_3 \in \mathcal{M}} f_1(x_1) + 2f_2(x_2) + f_3(x_3) + \rho(d_1(x_1, x_2)^2 + d_2(x_2, x_3)^2)/2$. To resolve this issue, we multiply the scalar α_i to f_i (and its Riemannian gradient) while updating x_is. Specifically, $\alpha_i = 1$ if $i = \{1, N\}$, else $\alpha_i = 0.5$. If $\mathrm{grad}_{x_i} f_i$ is the Riemannian gradient of f_i at $x_i \in \mathcal{M}$, then the stochastic gradient descent algorithm updates x_i along the search direction $-(\alpha_i \mathrm{grad}_{x_i} f_i + \rho \mathrm{grad}_{x_i} d_i)$ with the exponential mapping $\mathrm{Exp}_{x_i} : T_{x_i}\mathcal{M} \to \mathcal{M}$, where $T_{x_i}\mathcal{M}$ is the tangent space of $\mathcal{M}$ at x_i. The overall algorithm with concrete matrix expressions is in Table 1.

Table 2: Proposed preconditioned gossip algorithm for (7)

1. At each time slot t, pick an agent $i \leqslant N - 1$ randomly with uniform probability and compute the Riemannian gradients $\mathrm{grad}_{x_i} f_i$, $\mathrm{grad}_{x_{i+1}} f_{i+1}$, $\mathrm{grad}_{x_i} d_i$, and $\mathrm{grad}_{x_{i+1}} d_i$ with the matrix representations shown in Table 1.
2. Given a stepsize γ_t, update x_i and x_{i+1} as

$$x_{i+} = \mathrm{Exp}_{x_i}(-\gamma_t(\alpha_i \mathrm{grad}_{x_i} f_i + \rho \mathrm{grad}_{x_i} d_i)(\mathbf{W}_{i\mathbf{U}_i}^T \mathbf{W}_{i\mathbf{U}_i} + \rho \mathbf{I})^{-1})$$
$$x_{i+1+} = \mathrm{Exp}_{x_{i+1}}(-\gamma_t(\alpha_{i+1} \mathrm{grad}_{x_{i+1}} f_{i+1} + \rho \mathrm{grad}_{x_{i+1}} d_i)(\mathbf{W}_{i+1\mathbf{U}_{i+1}}^T \mathbf{W}_{i+1\mathbf{U}_{i+1}} + \rho \mathbf{I})^{-1}),$$

where $\mathbf{W}_{i\mathbf{U}_i}$ is the least-squares solution to the optimization problem $\min_{\mathbf{W}_i \in \mathbb{R}^{n_i \times r}} \|\mathcal{P}_{\Omega_i}(\mathbf{U}_i \mathbf{W}_i^T) - \mathcal{P}_{\Omega_i}(\mathbf{X}_i^\star)\|_F^2$. Exp and α_i are defined in Table 1.

The stochastic gradient descent algorithm in Table 1 converges to a critical point of (7) *almost surely* [20]. The gradient updates require the computation of the Riemannian gradient of the cost function in (7) and moving along the geodesics with *exponential* mapping, both of which have closed-form expressions on the Grassmann manifold $\mathrm{Gr}(r, m)$ [19]. Similarly, the matrix completion problem specific gradient computations are shown in [9].

3.1 Computational complexity

For an update of x_i with the formulas shown in Table 1, the computational complexity depends on the computation of *partial derivatives* of the cost function in (7), e.g., $\mathrm{Grad}_{x_i} f_i$. Particularly, in the context of the problem (5), the computational cost is $O(|\Omega_i|r^2 + n_i r^2 + mr)$. The Grassmann manifold related ingredients, e.g., Exp, cost $O(mr^2 + r^3)$.

3.2 Convergence analysis

Asymptotic convergence analysis of the algorithm in Table 1 follows directly from the analysis in [20, Theorem 1]. The key idea is that for a compact Riemannian manifold all continuous functions of the parameter can be bounded. This is the case for (7), which is on the *compact* Grassmann manifold $\mathrm{Gr}(r, m)$. Subsequently, under a decreasing stepsize condition and noisy gradient estimates (that is an unbiased estimator of the batch gradient), the stochastic gradient descent algorithm in Table 1 converges to a critical point of (7) almost surely. Conceptually, while the standard stochastic gradient descent setup deals with an *infinite* stream of samples, we deal with a finite number of samples (i.e., we pick an agent $i \leqslant N - 1$), which we repeat many times.

3.3 Preconditioned variant

The performance of first order algorithm (including stochastic gradients) often depends on the *condition number* (the ratio of maximum eigenvalue to the minimum eigenvalue) of the Hessian of the cost function (at the minimum). The issue of ill-conditioning arises especially when data $\mathbf{X}^\star$ have drawn power law distributed singular values. Additionally, a large value of ρ in (7) leads to convergence issues for numerical algorithms. To this end, the recent works [6, 7, 9] exploit the concept of *manifold preconditioning* in matrix completion. Specifically, the Riemannian gradients are *scaled* by computationally cheap matrix terms that arise from the second order curvature information of the cost function. Matrix scaling of the gradients is equivalent to multiplying an approximation of the *inverse* Hessian to gradients. This operation on a manifold requires special attention. In particular, the matrix scaling *must* be a positive definite operator on the tangent space of the manifold [7, 9].

Given the Riemannian gradient $\xi_{x_i} = \mathrm{grad}_{x_i} f_i + \rho \mathrm{grad}_{x_i} d_i$ computed by agent i, the proposed manifold preconditioning is

$$\xi_{x_i} \mapsto \xi_{x_i}(\underbrace{\mathbf{W}_{i\mathbf{U}_i}^T \mathbf{W}_{i\mathbf{U}_i}}_{\text{from completion}} + \underbrace{\rho \mathbf{I}}_{\text{from consensus}})^{-1}, \tag{8}$$

Table 3: Proposed parallel variant for (7)

---

1. Define round 1 as consisting of agents $i = 1, 3, \ldots$ and their neighbors. Define round 2 as consisting of agents $i = 2, 4, \ldots$ and their neighbors.
2. At each time slot t, pick a round $j \leqslant 2$ randomly with uniform probability.
3. Given a stepsize, update the agents (and their corresponding neighbors) in parallel with the updates proposed in Table 1 (or in Table 2).

---

where $\mathbf{W}_{i\mathbf{U}_i}$ is the solution to the optimization problem $\min_{\mathbf{W}_i \in \mathbb{R}^{n_i \times r}} \|\mathcal{P}_{\Omega_i}(\mathbf{U}_i \mathbf{W}_i^T) - \mathcal{P}_{\Omega_i}(\mathbf{X}_i^\star)\|_F^2$ and $\mathbf{I}$ is $r \times r$ identity matrix. The use of preconditioning (8) costs $O(n_i r^2 + r^3)$.

The term $\mathbf{W}_{i\mathbf{U}_i}^T \mathbf{W}_{i\mathbf{U}_i}$ is motivated by the fact that it is computationally cheap to compute and captures a *block diagonal approximation* of the Hessian of the simplified (but related) cost function $\|\mathbf{U}_i \mathbf{W}_{i\mathbf{U}_i}^T - \mathbf{X}_i^\star\|_F^2$. The works [6, 7, 9] use such *preconditioners* with superior performance. The term $\rho \mathbf{I}$ is motivated by the fact that the second order derivative of the square of the Riemannian geodesic distance is an identity matrix. Finally, it should be noted that the matrix scaling is positive definite, i.e., $\mathbf{W}_{i\mathbf{U}_i}^T \mathbf{W}_{i\mathbf{U}_i} + \rho \mathbf{I} \succ 0$ and that the transformation (8) is on the tangent space. Equivalently, if ξ_{x_i} belongs to $T_{x_i}\mathcal{M}$, then $\xi_{x_i}(\mathbf{W}_{i\mathbf{U}_i}^T \mathbf{W}_{i\mathbf{U}_i} + \rho \mathbf{I})^{-1}$ also belongs to $T_{x_i}\mathcal{M}$. This is readily checked by the fact that the tangent space $T_{x_i}\mathcal{M}$ at x_i on the Grassmann manifold is characterized by the set $\{\eta_{x_i} : \eta_{x_i} \in \mathbb{R}^{m \times r}, \mathbf{U}_i^T \eta_{x_i} = 0\}$.

The proposed preconditioned variant of the stochastic gradient descent algorithm for (7) is shown in Table 2. It should be noted that preconditioning the gradients does not affect the asymptotic convergence guarantees of the proposed algorithm.

3.4 Parallel variant

Assumption **A1** on the network topology of agents allows to propose parallel variants of the proposed stochastic gradient descent algorithms in Tables 1 and 2. To this end, instead of picking one agent at a time, we pick agents in such a way that it leads to a number of parallel updates.

We explain the idea for $N = 5$. Updates of the agents are divided into two *rounds*. In round 1, we pick agents 1 and 3, i.e., all the *odd* numbered agents. It should be noted that the neighbor of agent 1 is agent 2 and the neighbor of agent 3 is agent 4. Consequently, the updates of agents 1 and 2 are independent from those of agents 3 and 4 and hence, can be carried out in parallel. In round 2, we pick agents 2 and 4, i.e., all the *even* numbered agents. The updates of agents 2 and 3 are independent from those of agents 4 and 5 and therefore, can be carried out in parallel.

The key idea is that *randomness* is on the rounds and not on the agents. For example, we pick a round j from $\{1, 2\}$ with uniform probability. Once a round is picked, the updates on the agents (that are part of this round) are performed with the same stepsize and in parallel. The stepsize is updated when a new round is picked. The stepsize sequence satisfies the standard conditions, i.e., it is square-summable and its summation is divergent. The overall algorithm in shown in Table 3.

To prove convergence, we define two new functions,

$$
\begin{aligned}
g_1(x_1, x_2, \ldots) &= f_1(x_1) + f_2(x_2) + \ldots + \frac{\rho}{2}(d_1(x_1, x_2)^2 + d_3(x_3, x_4)^2 + \ldots) \\
g_2(x_2, x_3, \ldots) &= f_2(x_2) + f_3(x_3) + \ldots + \frac{\rho}{2}(d_2(x_2, x_3)^2 + d_4(x_4, x_5)^2 + \ldots),
\end{aligned}
\tag{9}
$$

that consist of terms from the cost function in (7). The algorithm in Table 3 is then interpreted as the standard stochastic gradient descent algorithm applied to the problem

$$
\min_{x_i \in \mathcal{M}} \quad g_1(x_1, x_2, \ldots) + g_2(x_2, x_3, \ldots).
\tag{10}
$$

with two "samples" that are chosen randomly at each time slot. Consequently, following the standard arguments, the algorithm in Table 3 converges asymptotically to a critical point of (10). However, it should also be noted that the addition of g_1 and g_2 leads to x_2 to x_{N-1} being updated (on an average) *twice* the number of times x_1 and x_N are updated. This is handled by multiplying α_i to

Table 4: Proposed algorithm for continuously changing network topology

1. At each time slot t, pick a pair of agents, say i and k, randomly with uniform probability.
2. Compute the Riemannian gradients $\text{grad}_{x_i} f_i$, $\text{grad}_{x_k} f_k$, $\text{grad}_{x_i} d_{ik}$, and $\text{grad}_{x_k} d_{ik}$ as

$$
\begin{aligned}
\text{Grad}_{x_i} f_i &= (\mathcal{P}_{\Omega_i}(\mathbf{U}_i \mathbf{W}_{i\mathbf{U}_i}^T) - \mathcal{P}_{\Omega_i}(\mathbf{X}_i^\star))\mathbf{W}_{i\mathbf{U}_i} \\
\text{grad}_{x_i} f_i &= \text{Grad}_{x_i} f_i - \mathbf{U}_i(\mathbf{U}_i^T \text{Grad}_{x_i} f_i) \\
\text{grad}_{x_i} d_{ik} &= -\text{Log}_{x_i}(x_k) \\
\text{grad}_{x_k} d_{ik} &= -\text{Log}_{x_k}(x_i),
\end{aligned}
$$

where Log is defined in Table 1.
3. Given a stepsize γ_t, update x_i and x_k as

$$
\begin{aligned}
x_{i+} &= \text{Exp}_{x_i}(-\gamma_t(\text{grad}_{x_i} f_i + \rho\,\text{grad}_{x_i} d_{ik})) \\
x_{k+} &= \text{Exp}_{x_k}(-\gamma_t(\text{grad}_{x_k} f_k + \rho\,\text{grad}_{x_k} d_{ik})),
\end{aligned}
$$

where the exponential mapping Exp_{x_i} is defined in Table 1.

f_i while updating x_is, where $\alpha_i = 1$ if $i = \{1, N\}$, else $\alpha_i = 0.5$. Finally, the algorithm in Table (3) converges to a critical point of (7). It should emphasized that parallelization of the updates is for free by virtue of construction of functions in (9).

3.5 Extension to continuously changing network topology

The algorithm in Table 1 assumes that the neighbors of the agents are predefined in a particular way (assumption **A1**). However, in many scenarios the network topology changes with time [16]. To simulate the scenario, we first consider a fully connected network of N agents. The number of unique edges is $N(N-1)/2$. We pick an edge ik (the edge that connects agents i and k) randomly with uniform probability and drop all the other edges. Equivalently, only one edge is *active* at any time slot. Consequently, we update agents i and k with a gradient descent update, e.g., based on Table 1 or Table 2. The overall algorithm is shown in Table 4. Following the arguments in Section 3.2, it is straightforward to see that the proposed algorithm converges almost surely to a critical point of a problem that combines completion along with consensus, i.e.,

$$
\min_{x_1,\ldots,x_N \in \mathcal{M}} \quad (N-1)\sum_i f_i(x_i) + \frac{\rho}{2}\sum_{i<k} d_{ik}(x_i, x_k)^2, \tag{11}
$$

where $d_{ik}(x_i, x_k)$ is the Riemannian geodesic distance between x_i and x_k.

4 Numerical comparisons

Our proposed algorithms in Table 1 (Online Gossip) and in Table 2 (Precon Online Gossip) and their parallel variants, Parallel Gossip and Precon Parallel Gossip, are compared on different problem instances. The implementations are based on the Manopt toolbox [22] with certain operations relying on the mex files supplied with [9]. We also show comparisons with D-LMaFit, the decentralized algorithm proposed in [15] on smaller instances as the D-LMaFit code (supplied by the authors) is *not* tuned to large-scale instances. As the mentioned algorithms are well suited for different scenarios, we compare them against the number of *updates* performed by the agents. We fix the number of agents N to 6. Online algorithms are run for a maximum of 1000 iterations. The parallel variants are run for 400 iterations. Overall, agents 1 and N perform a maximum of 200 updates (rest all perform 400 updates). D-LMaFit is run for 400 iterations, i.e., each agent performs 400 updates. Algorithms are initialized randomly. The stepsize sequence is defined as $\gamma_t = \gamma_0/t$, where t is the time slot and γ_0 is set using cross validation. For simplicity, all figures only show the plots for agents 1 and 2.

All simulations are performed in Matlab on a 2.7 GHz Intel Core i5 machine with 8 GB of RAM. For each example considered here, an $m \times n$ random matrix of rank r is generated as in [4]. Two matrices $\mathbf{A} \in \mathbb{R}^{m \times r}$ and $\mathbf{B} \in \mathbb{R}^{n \times r}$ are generated according to a Gaussian distribution with zero

mean and unit standard deviation. The matrix product $\mathbf{AB}^T$ gives a random matrix of rank r. A fraction of the entries are randomly removed with uniform probability and noise (sampled from the Gaussian distribution with mean zero and standard deviation 10^{-6}) is added to each entry to construct the training set Ω and $\mathbf{X}^\star$. The over-sampling ratio (OS) is the ratio of the number of known entries to the matrix dimension, i.e, $\text{OS} = |\Omega|/(mr + nr - r^2)$. We also create a test set by randomly picking a small set of entries from $\mathbf{AB}^T$. The matrices $\mathbf{X}_i^\star$ are created by distributing the number of n columns of $\mathbf{X}^\star$ equally among the agents. The training and test sets are also partitioned similarly.
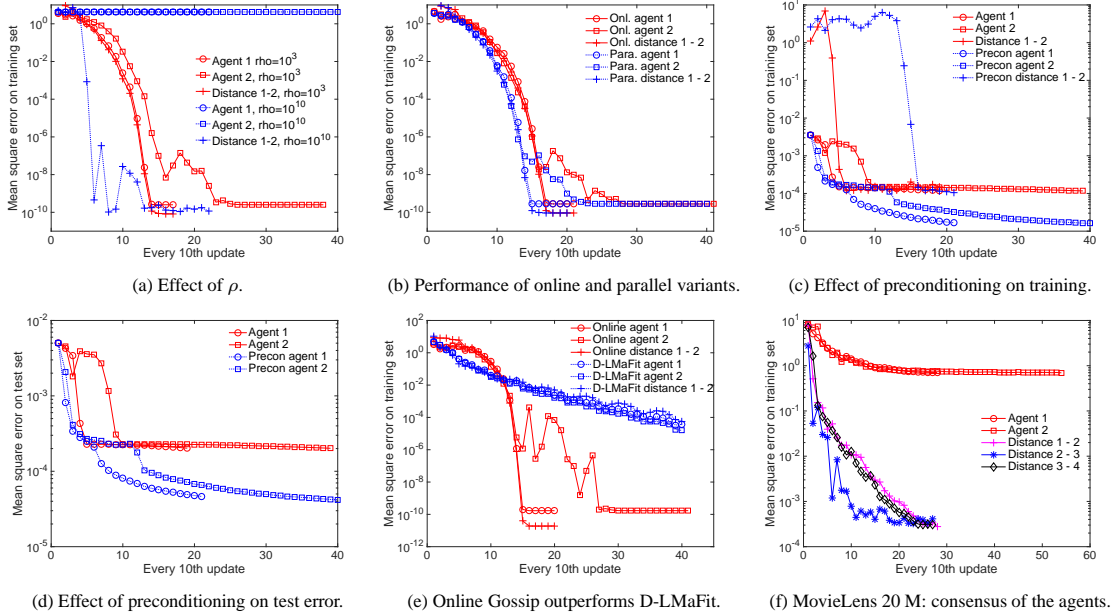

(a) Effect of ρ.

(b) Performance of online and parallel variants.

(c) Effect of preconditioning on training.

(d) Effect of preconditioning on test error.

(e) Online Gossip outperforms D-LMaFit.

(f) MovieLens 20 M: consensus of the agents.

Figure 1: Performance of proposed algorithms in different scenarios.

Case 1: effect of ρ. We consider a problem instance of size $10\,000 \times 100\,000$ of rank 5 and OS 6. Two scenarios with $\rho = 10^3$ and $\rho = 10^{10}$ are considered. Figure 1(a) shows the performance of Online Gossip. Not surprisingly, for $\rho = 10^{10}$, we only see consensus (the distance between agents 1 and 2 tends to zero). For $\rho = 10^3$, we see both completion and consensus, which validates the theory.

Case 2: performance of online versus parallel. We consider Case 1 with $\rho = 10^3$. Figure 1(b) shows the performance of Online Gossip and Parallel Gossip, both of which show a similar behavior on the training and test (not shown here) sets.

Case 3: ill-conditioned instances. We consider a problem instance of size $5\,000 \times 50\,000$ of rank 5 and impose an exponential decay of singular values with condition number 500 and OS 6. Figure 1(c) shows the performance of Online Gossip and its preconditioned variant for $\rho = 10^3$. During the initial updates, the preconditioned variant aggressively minimizes the completion term of (5), which shows the effect of the preconditioner (8). Eventually, consensus among the agents is achieved. Overall, the preconditioned variant shows a superior performance on both the training and test sets as shown in Figures 1(c) and 1(d).

Case 4: Comparisons with D-LMaFit [15]. We consider a problem instance of size 500×12000, rank 5, and OS 6. D-LMaFit is run with the default parameters. For Online Gossip, we set $\rho = 10^3$. As shown in Figure 1(e), Online Gossip quickly outperforms D-LMaFit. Overall, Online Gossip takes fewer number of updates to reach a high accuracy.

Case 5: MovieLens 20M dataset [23]. Finally, we show the performance of Online Gossip on the MovieLens-20M dataset of 20000263 ratings by 138493 users for 26744 movies. (D-LMaFit is not compared as it does not scale to this dataset.) We perform 5 random 80/20 train/test partitions. We split both the train and test data among $N = 4$ agents along the number of users such that

each agent has ratings for 26744 movies and 34624 (except agent 4, which has 34621) unique users. This ensures that the ratings are distributed evenly among the agents. We run Online Gossip with $\rho = 10^7$ (through cross validation) and for 800 iterations. Figure 1(f) shows that asymptotic consensus is achieved among the four agents. It should be noted that the distance between agents 2 and 3 decreases faster than others as agents 2 and 3 are updated (on an average) twice the number of times than agents 1 and 4 (assumption **A1**). Table 5 shows the normalized mean absolute errors (NMAE) obtained on the *full* test set averaged over five runs. NMAE is defined as the mean absolute error (MAE) divided by variation of the ratings. Since the ratings vary from 0.5 to 5, NMAE is MAE/4.5. We obtain the lowest NMAE for rank 5.

Table 5: Performance of Online Gossip on MovieLens 20M dataset

	Rank 3	Rank 5	Rank 7	Rank 9
NMAE on test set	$0.1519 \pm 3 \cdot 10^{-3}$	$\mathbf{0.1507 \pm 3 \cdot 10^{-3}}$	$0.1531 \pm 2 \cdot 10^{-3}$	$0.1543 \pm 1 \cdot 10^{-3}$

5   Conclusion

We have proposed a Riemannian gossip approach to the decentralized matrix completion problem. Specifically, the completion task is distributed among a number of agents, which are then required to achieve consensus. Exploiting the gossip framework, this is modeled as minimizing a weighted sum of *completion* and *consensus* terms on the Grassmann manifold. The rich geometry of the Grassmann manifold allowed to propose a novel stochastic gradient descent algorithm for the problem with simple updates. Additionally, we have proposed two variants – preconditioned and parallel – of the algorithm for dealing with different scenarios. Numerical experiments show the competitive performance of the proposed algorithms on different benchmarks.

References

[1] I. Markovsky and K. Usevich. Structured low-rank approximation with missing data. *SIAM Journal on Matrix Analysis and Applications*, 34(2):814–830, 2013.

[2] J. Rennie and N. Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *International Conference on Machine learning (ICML)*, pages 713–719, 2005.

[3] J. Shi, Y.and Zhang and K. B. Letaief. Low-rank matrix completion for topological interference management by Riemannian pursuit. *IEEE Transactions on Wireless Communications*, PP(99), 2016.

[4] J. F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.

[5] K. C. Toh and S. Yun. An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems. *Pacific Journal of Optimization*, 6(3):615–640, 2010.

[6] T. T. Ngo and Y. Saad. Scaled gradients on Grassmann manifolds for matrix completion. In *Advances in Neural Information Processing Systems 25 (NIPS)*, pages 1421–1429, 2012.

[7] B. Mishra and R. Sepulchre. R3MC: A Riemannian three-factor algorithm for low-rank matrix completion. In *Proceedings of the 53rd IEEE Conference on Decision and Control (CDC)*, pages 1137–1142, 2014.

[8] Z. Wen, W. Yin, and Y. Zhang. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Mathematical Programming Computation*, 4(4):333–361, 2012.

[9] N. Boumal and P.-A. Absil. Low-rank matrix completion via preconditioned optimization on the Grassmann manifold. *Linear Algebra and its Applications*, 475:200–239, 2015.

[10] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. *IEEE Transactions on Information Theory*, 56(6):2980–2998, 2010.

[11] L. Balzano, R. Nowak, and B. Recht. Online identification and tracking of subspaces from highly incomplete information. In *The 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 704–711, June 2010.

[12] H.-F. Yu, C.-J. Hsieh, S. Si, and I. S. Dhillon. Parallel matrix factorization for recommender systems. *Knowledge and Information Systems*, 41(3):793–819, 2014.

[13] B. Recht and C Ré. Parallel stochastic gradient algorithms for large-scale matrix completion. *Mathematical Programming Computation*, 5(2):201–226, 2013.

[14] C. Teflioudi, F. Makari, and R. Gemulla. Distributed matrix completion. In *International Conference on Data Mining (ICDM)*, pages 655–664, 2012.

[15] A.-Y. Lin and Q. Ling. Decentralized and privacy-preserving low-rank matrix completion. *Journal of the Operations Research Society of China*, 3(2):189–205, 2015.

[16] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transaction on Information Theory*, 52(6):2508–2530, 2006.

[17] B. Mishra, G. Meyer, S. Bonnabel, and R. Sepulchre. Fixed-rank matrix factorizations and Riemannian low-rank optimization. *Computational Statistics*, 29(3–4):591–621, 2014.

[18] W. Dai, E. Kerman, and O. Milenkovic. A geometric approach to low-rank matrix completion. *IEEE Transactions on Information Theory*, 58(1):237–247, 2012.

[19] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008.

[20] S. Bonnabel. Stochastic gradient descent on Riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229, 2013.

[21] K. Hüper, U. Helmke, and S. Herzberg. On the computation of means on Grassmann manifolds. In *International Symposium on Mathematical Theory of Networks and Systems (MTNS)*, pages 2439–2441, 2010.

[22] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt: a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15(Apr):1455–1459, 2014.

[23] F. M. Harper and J. A. Konstan. The MovieLens datasets: history and contex. *ACM Transactions on Interactive Intelligent Systems*, 5(4), 2015.