

AN APOCALYPSE-FREE FIRST-ORDER LOW-RANK OPTIMIZATION ALGORITHM WITH AT MOST ONE RANK REDUCTION ATTEMPT PER ITERATION*

GUILLAUME OLIKIER[†] AND P.-A. ABSIL[†]

Abstract. We consider the problem of minimizing a differentiable function with locally Lipschitz continuous gradient on the real determinantal variety and present a first-order algorithm designed to find a stationary point of that problem. This algorithm applies steps of a retraction-free descent method proposed by Schneider and Uschmajew [*SIAM J. Optim.*, 25 (2015), pp. 622–646], while taking the numerical rank into account to attempt rank reductions. We prove that this algorithm produces a sequence of iterates whose accumulation points are stationary and therefore does not follow the so-called apocalypses described by Levin, Kileel, and Boumal [*Math. Program.*, 199 (2023), pp. 831–864]. Moreover, the rank reduction mechanism of this algorithm requires at most one rank reduction attempt per iteration, in contrast with the one of the P²GDR algorithm introduced by Olikier, Gallivan, and Absil [*An Apocalypse-Free First-Order Low-Rank Optimization Algorithm*, Technical report UCL-INMA-2022.01, <https://arxiv.org/abs/2201.03962>, 2022], which can require a number of rank reduction attempts equal to the rank of the iterate in the worst-case scenario.

Key words. convergence analysis, stationarity, low-rank optimization, determinantal variety, first-order method, tangent cones, QR factorization, singular value decomposition

MSC codes. 14M12, 65K10, 90C26, 90C30, 40A05

DOI. 10.1137/22M1518256

1. Introduction. As in [7], we consider the problem

$$(1) \quad \min_{X \in \mathbb{R}_{\leq r}^{m \times n}} f(X)$$

of minimizing a differentiable function $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ with locally Lipschitz continuous gradient on the *determinantal variety* [4, Lecture 9]

$$(2) \quad \mathbb{R}_{\leq r}^{m \times n} := \{X \in \mathbb{R}^{m \times n} \mid \text{rank } X \leq r\},$$

m , n , and r being positive integers such that $r < \min\{m, n\}$. This problem appears in several applications such as matrix equations, model reduction, matrix sensing, and matrix completion; see, e.g., [11, 3] and the references therein. As problem (1) is, in general, intractable [1], our goal is to find a *stationary point* of this problem, i.e., a zero of the *stationarity measure*

$$(3) \quad s(\cdot; f, \mathbb{R}_{\leq r}^{m \times n}) : \mathbb{R}_{\leq r}^{m \times n} \rightarrow \mathbb{R} : X \mapsto \|P_{T_{\mathbb{R}_{\leq r}^{m \times n}}(X)}(-\nabla f(X))\|$$

that returns the norm of any projection of $-\nabla f(X)$ onto the tangent cone to $\mathbb{R}_{\leq r}^{m \times n}$ at X ; the notation is introduced in section 2. Indeed, by [5, Lemmas A.7 and A.8], the correspondence

*Received by the editors August 25, 2022; accepted for publication (in revised form) by D. Orban April 3, 2023; published electronically September 22, 2023.

<https://doi.org/10.1137/22M1518256>

Funding: This work was supported by the Fonds de la Recherche Scientifique (FNRS) and the Fonds Wetenschappelijk Onderzoek – Vlaanderen under EOS Project 30468160.

[†]ICTEAM Institute, UCLouvain, Avenue Georges Lemaître 4, 1348 Louvain-la-Neuve, Belgium (guillaume.olikier@uclouvain.be, pa.absil@uclouvain.be).

$$\mathbb{R}_{\leq r}^{m \times n} \dashrightarrow \mathbb{R}^{m \times n} : X \mapsto P_{T_{\mathbb{R}_{\leq r}^{m \times n}}(X)}(-\nabla f(X))$$

depends on f only through its restriction $f|_{\mathbb{R}_{\leq r}^{m \times n}}$, and, by [10, Theorem 6.12] and [5, Proposition 2.5], if $X \in \mathbb{R}_{\leq r}^{m \times n}$ is a local minimizer of $f|_{\mathbb{R}_{\leq r}^{m \times n}}$, then $s(X; f, \mathbb{R}_{\leq r}^{m \times n}) = 0$.

To the best of our knowledge, the second-order method given in [5, Algorithm 1] and the first-order method given in [7, Algorithm 2], dubbed P²GDR, are the only two algorithms in the literature that provably accumulate at stationary points of (1). Other algorithms, such as [11, Algorithm 3], known as P²GD, and [11, Algorithm 4], which we call RFD because it is a retraction-free descent method, can fail as they can produce a feasible sequence $(X_i)_{i \in \mathbb{N}}$ that converges to some point X with the property that $\lim_{i \rightarrow \infty} s(X_i; f, \mathbb{R}_{\leq r}^{m \times n}) = 0 < s(X; f, \mathbb{R}_{\leq r}^{m \times n})$. Such a triplet $(X, (X_i)_{i \in \mathbb{N}}, f)$ is called an *apocalypse* and the point X , which necessarily satisfies $\text{rank } X < r$, is said to be *apocalyptic* [5, Definition 2.7].

P²GDR is P²GD equipped with a rank reduction mechanism. On the one hand, this mechanism ensures the apocalypse-free property. On the other hand, as mentioned in [7, section 6], it can generate a significant computational overhead in some situations. Indeed, its for loop can consider up to r rank reductions of the current iterate to each of which the P²GD map [7, Algorithm 1] is applied. This is further discussed in section 7.

In this paper, we introduce a first-order optimization algorithm on $\mathbb{R}_{\leq r}^{m \times n}$ (Algorithm 3), called RFDR, that accumulates at stationary points (Theorem 6.2) and is thus apocalypse-free. We also deduce that $s(\cdot; f, \mathbb{R}_{\leq r}^{m \times n})$ goes to zero along every convergent subsequence of the generated sequence (Corollary 6.3). RFDR is RFD equipped with a rank reduction mechanism. RFDR and P²GDR, which is the only other first-order optimization algorithm on $\mathbb{R}_{\leq r}^{m \times n}$ known to accumulate at stationary points, have their own advantages. On the one hand, the search directions used by P²GD are more closely related to the negative gradient than those used by RFD; while this does not imply that P²GD converges faster than RFD, such an observation was made experimentally in [11, section 3.4] on a matrix completion problem. On the other hand, RFDR presents two advantages over P²GDR. First, it inherits the advantage that RFD has over P²GD, namely being retraction-free; in other words, the updates are performed along a straight line. Second, its rank reduction mechanism is more efficient than the one of P²GDR because it ensures the apocalypse-free property by performing at most one rank reduction attempt per iteration (compare Algorithm 2 with [7, Algorithm 3]).

This paper is organized as follows. After introducing the notation in section 2, we define in section 3 the *restricted tangent cone* to $\mathbb{R}_{\leq r}^{m \times n}$; the descent direction used by RFD is the projection of the negative gradient onto that closed cone. In section 4, we analyze the iteration map of RFD (Algorithm 1) under the assumption of local Lipschitz continuity of ∇f . We introduce the RFDR algorithm in section 5, analyze its convergence properties in section 6, and compare its computational cost with the one of P²GDR in section 7. Section 8 contains concluding remarks.

2. Notation. In this section, we introduce the notation used throughout the paper. This is the same notation as in [7], to which we refer for a more complete review of the background material. The real vector space $\mathbb{R}^{m \times n}$ is endowed with the Frobenius inner product $\langle \cdot, \cdot \rangle$, $\|\cdot\|$ denotes the Frobenius norm, and, for all $X \in \mathbb{R}^{m \times n}$ and all $\rho \in (0, \infty)$, $B(X, \rho)$ and $B[X, \rho]$ respectively denote the open ball and the closed ball of center X and radius ρ in $\mathbb{R}^{m \times n}$. Given a nonempty subset \mathcal{S} of $\mathbb{R}^{m \times n}$, the tangent cone to \mathcal{S} at $X \in \mathcal{S}$ is denoted by $T_{\mathcal{S}}(X)$, the distance from $X \in \mathbb{R}^{m \times n}$

to \mathcal{S} is denoted by $d(X, \mathcal{S})$, and the projection of $X \in \mathbb{R}^{m \times n}$ onto \mathcal{S} is denoted by $P_{\mathcal{S}}(X)$. For every $r \in \{0, \dots, \min\{m, n\}\}$,

$$\mathbb{R}_r^{m \times n} := \{X \in \mathbb{R}^{m \times n} \mid \text{rank } X = r\}$$

is the smooth manifold of $m \times n$ rank- r real matrices, and, if $r \geq 1$,

$$\text{St}(r, n) := \{U \in \mathbb{R}^{n \times r} \mid U^T U = I_r\}$$

is a Stiefel manifold. We also write $\mathbb{R}_*^{m \times n} := \mathbb{R}_{\min\{m, n\}}^{m \times n}$. The singular values of $X \in \mathbb{R}^{m \times n}$ are denoted by $\sigma_1(X) \geq \dots \geq \sigma_{\min\{m, n\}}(X) \geq 0$.

3. The restricted tangent cone to the determinantal variety. In this section, we define the restricted tangent cone to $\mathbb{R}_{\leq r}^{m \times n}$ (Definition 3.1) and prove that the descent direction introduced in [11, section 3.4] is the projection of the negative gradient onto that closed cone (Proposition 3.2). Throughout this section, $r \in \{1, \dots, r\}$, $X \in \mathbb{R}_r^{m \times n}$, $U \in \text{St}(r, m)$, $U_{\perp} \in \text{St}(m - r, m)$, $V \in \text{St}(r, n)$, $V_{\perp} \in \text{St}(n - r, n)$, $\text{im } U = \text{im } X$, $\text{im } U_{\perp} = (\text{im } X)^{\perp}$, $\text{im } V = \text{im } X^T$, and $\text{im } V_{\perp} = (\text{im } X^T)^{\perp}$.

The set introduced in the following definition is a closed cone contained in

$$T_{\mathbb{R}_{\leq r}^{m \times n}}(X) = \left\{ [U \ U_{\perp}] \begin{bmatrix} A & B \\ C & D \end{bmatrix} [V \ V_{\perp}]^T \mid \begin{array}{l} A \in \mathbb{R}^{r \times r}, B \in \mathbb{R}^{r \times n-r}, \\ C \in \mathbb{R}^{m-r \times r}, D \in \mathbb{R}_{\leq r-r}^{m-r \times n-r} \end{array} \right\}.$$

DEFINITION 3.1. *The restricted tangent cone to $\mathbb{R}_{\leq r}^{m \times n}$ at X is*

$$\check{T}_{\mathbb{R}_{\leq r}^{m \times n}}(X) := \left\{ [U \ U_{\perp}] \begin{bmatrix} A & B \\ C & D \end{bmatrix} [V \ V_{\perp}]^T \mid \begin{array}{l} A \in \mathbb{R}^{r \times r}, B \in \mathbb{R}^{r \times n-r}, \\ C \in \mathbb{R}^{m-r \times r}, D \in \mathbb{R}_{\leq r-r}^{m-r \times n-r}, \\ B = 0_{r \times n-r} \text{ or } C = 0_{m-r \times r} \end{array} \right\}.$$

Furthermore,

$$\check{T}_{\mathbb{R}_{\leq r}^{m \times n}}(0_{m \times n}) := T_{\mathbb{R}_{\leq r}^{m \times n}}(0_{m \times n}) = \mathbb{R}_{\leq r}^{m \times n}.$$

The following proposition can be compared to [7, Proposition 2.7].

PROPOSITION 3.2. *Let $Z \in \mathbb{R}^{m \times n}$ be written as*

$$Z = [U \ U_{\perp}] \begin{bmatrix} A & B \\ C & D \end{bmatrix} [V \ V_{\perp}]^T$$

with $A = U^T Z V$, $B = U^T Z V_{\perp}$, $C = U_{\perp}^T Z V$, and $D = U_{\perp}^T Z V_{\perp}$. Then,

$$(4) \quad P_{\check{T}_{\mathbb{R}_{\leq r}^{m \times n}}(X)}(Z) = \begin{cases} [U \ U_{\perp}] \begin{bmatrix} A & B \\ 0_{m-r \times r} & P_{\mathbb{R}_{\leq r-r}^{m-r \times n-r}}(D) \end{bmatrix} [V \ V_{\perp}]^T & \text{if } \|B\| > \|C\|, \\ [U \ U_{\perp}] \begin{bmatrix} A & B \\ 0_{m-r \times r} & P_{\mathbb{R}_{\leq r-r}^{m-r \times n-r}}(D) \end{bmatrix} [V \ V_{\perp}]^T & \text{if } \|B\| = \|C\|, \\ \cup [U \ U_{\perp}] \begin{bmatrix} A & 0_{r \times n-r} \\ C & P_{\mathbb{R}_{\leq r-r}^{m-r \times n-r}}(D) \end{bmatrix} [V \ V_{\perp}]^T & \\ [U \ U_{\perp}] \begin{bmatrix} A & 0_{r \times n-r} \\ C & P_{\mathbb{R}_{\leq r-r}^{m-r \times n-r}}(D) \end{bmatrix} [V \ V_{\perp}]^T & \text{if } \|B\| < \|C\|. \end{cases}$$

Moreover,

$$(5) \quad \left\| P_{T_{\mathbb{R}_{\leq r}^{m \times n}}(X)}(Z) \right\| \geq \left\| P_{\check{T}_{\mathbb{R}_{\leq r}^{m \times n}}(X)}(Z) \right\| \geq \frac{1}{\sqrt{2}} \left\| P_{T_{\mathbb{R}_{\leq r}^{m \times n}}(X)}(Z) \right\|$$

and, for all $\check{Z} \in P_{\check{T}_{\mathbb{R}_{\leq r}^{m \times n}}(X)}(Z)$,

$$(6) \quad \langle Z, \check{Z} \rangle = \|\check{Z}\|^2.$$

Furthermore,

$$(7) \quad \|Z\| \geq \left\| P_{T_{\mathbb{R}_{\leq r}^{m \times n}}(X)}(Z) \right\| \geq \sqrt{\frac{r-r_l}{\min\{m,n\}-r_l}} \|Z\|.$$

Proof. We prove only (4) as (5) and (6) are given in [11, section 3.4] and (7) in [11, (3.7)]. Let $\check{Z} \in \check{T}_{\mathbb{R}_{\leq r}^{m \times n}}(X)$. Then,

$$\check{Z} = [U \ U_{\perp}] \begin{bmatrix} \check{A} & \check{B} \\ \check{C} & \check{D} \end{bmatrix} [V \ V_{\perp}]^{\top}$$

for some $\check{A} \in \mathbb{R}^{r \times r}$, $\check{B} \in \mathbb{R}^{r \times n-r}$, $\check{C} \in \mathbb{R}^{m-r \times r}$, and $\check{D} \in \mathbb{R}_{\leq r-r}^{m-r \times n-r}$ such that $\check{B} = 0_{r \times n-r}$ or $\check{C} = 0_{m-r \times r}$. Thus,

$$\begin{aligned} \|Z - \check{Z}\|^2 &= \left\| \begin{bmatrix} A - \check{A} & B - \check{B} \\ C - \check{C} & D - \check{D} \end{bmatrix} \right\|^2 \\ &= \|A - \check{A}\|^2 + \|B - \check{B}\|^2 + \|C - \check{C}\|^2 + \|D - \check{D}\|^2 \\ &\geq 0 + \min\{\|B\|, \|C\|\}^2 + d(D, \mathbb{R}_{\leq r-r}^{m-r \times n-r})^2. \end{aligned}$$

Therefore,

$$d(Z, \check{T}_{\mathbb{R}_{\leq r}^{m \times n}}(X))^2 \geq \min\{\|B\|, \|C\|\}^2 + d(D, \mathbb{R}_{\leq r-r}^{m-r \times n-r})^2,$$

a bound which is reached only by (4). □

In practice, the projection onto $T_{\mathbb{R}_{\leq r}^{m \times n}}(X)$ can be computed thanks to [11, Algorithm 2]. This does not rely on U_{\perp} and V_{\perp} , which are huge in the frequently encountered case where $r \ll \min\{m, n\}$. The practical computation of the projection onto $\check{T}_{\mathbb{R}_{\leq r}^{m \times n}}(X)$ is described in [11, section 3.4] and does not rely on U_{\perp} and V_{\perp} either. We discuss the practical implementation and the computational cost in section 7.

4. The RFD map. In this section, we analyze Algorithm 1—which corresponds to the iteration map of RFD [11, Algorithm 4] except that the initial step size for the backtracking procedure is chosen in a given bounded interval—under the assumption that f is differentiable with ∇f locally Lipschitz continuous. This serves as a basis for the convergence analysis conducted in section 6 since the RFD map (Algorithm 1) is used as a subroutine by the RFDR map (Algorithm 2). The analysis conducted in this section is to RFD as the analysis conducted in [7, section 3] is to P²GD, although it is simpler since RFD does not require any retraction as $X + \alpha G \in \mathbb{R}_{\leq r}^{m \times n}$ for all $X \in \mathbb{R}_{\leq r}^{m \times n}$, all $\alpha \in (0, \infty)$, and all $G \in \check{T}_{\mathbb{R}_{\leq r}^{m \times n}}(X)$.

Let us recall that, since ∇f is locally Lipschitz continuous, for every closed ball $\mathcal{B} \subset \mathbb{R}^{m \times n}$,

$$\text{Lip}(\nabla f) := \sup_{\substack{X, Y \in \mathcal{B} \\ X \neq Y}} \frac{\|\nabla f(X) - \nabla f(Y)\|}{\|X - Y\|} < \infty,$$

Algorithm 1. RFD map (based on [11, Algorithm 4]).

Require: $(f, r, \underline{\alpha}, \bar{\alpha}, \beta, c)$ where $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is differentiable with ∇f locally Lipschitz continuous, $r < \min\{m, n\}$ is a positive integer, $0 < \underline{\alpha} \leq \bar{\alpha} < \infty$, and $\beta, c \in (0, 1)$.

Input: $X \in \mathbb{R}_{\leq r}^{m \times n}$ such that $s(X; f, \mathbb{R}_{\leq r}^{m \times n}) > 0$.

Output: a point in $\text{RFD}(X; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c)$.

1: Choose $G \in P_{\mathbb{R}_{\leq r}^{m \times n}(X)}(-\nabla f(X))$ and $\alpha \in [\underline{\alpha}, \bar{\alpha}]$;

2: **while** $f(X + \alpha G) > f(X) - c\alpha\|G\|^2$ **do**

3: $\alpha \leftarrow \alpha\beta$;

4: **end while**

5: Return $X + \alpha G$.

which implies, by [6, Lemma 1.2.3], that, for all $X, Y \in \mathcal{B}$,

$$(8) \quad |f(Y) - f(X) - \langle \nabla f(X), Y - X \rangle| \leq \frac{\text{Lip}_{\mathcal{B}}(\nabla f)}{2} \|Y - X\|^2.$$

PROPOSITION 4.1. *Let $X \in \mathbb{R}_{\leq r}^{m \times n}$ and $\bar{\alpha} \in (0, \infty)$. Let $\mathcal{B} \subsetneq \mathbb{R}^{m \times n}$ be a closed ball such that, for all $G \in P_{\mathbb{R}_{\leq r}^{m \times n}(X)}(-\nabla f(X))$ and all $\alpha \in [0, \bar{\alpha}]$, $X + \alpha G \in \mathcal{B}$; an example of such a ball is $B[X, \bar{\alpha}s(X; f, \mathbb{R}_{\leq r}^{m \times n})]$. Then, for all $G \in P_{\mathbb{R}_{\leq r}^{m \times n}(X)}(-\nabla f(X))$ and all $\alpha \in [0, \bar{\alpha}]$,*

$$(9) \quad f(X + \alpha G) \leq f(X) + \|G\|^2 \alpha \left(-1 + \frac{\text{Lip}_{\mathcal{B}}(\nabla f)}{2} \alpha \right).$$

Proof. The example $B[X, \bar{\alpha}s(X; f, \mathbb{R}_{\leq r}^{m \times n})]$ is correct by the first inequality of (5). The inequality (9) is based on (8) and (6):

$$\begin{aligned} f(X + \alpha G) - f(X) &\leq \langle \nabla f(X), (X + \alpha G) - X \rangle + \frac{\text{Lip}_{\mathcal{B}}(\nabla f)}{2} \|(X + \alpha G) - X\|^2 \\ &= -\alpha\|G\|^2 + \frac{\text{Lip}_{\mathcal{B}}(\nabla f)}{2} \alpha^2 \|G\|^2. \quad \square \end{aligned}$$

Let us make two remarks concerning the preceding proposition. First, the existence of a ball \mathcal{B} crucially relies on the upper bound $\bar{\alpha}$ required by Algorithm 1. Second, in contrast with [7, equation (8)], the upper bound (9) does not depend on the curvature of a fixed-rank manifold. In particular, (9) does not involve any singular value. This fundamental difference is the reason why RFDR is apocalypse-free while requiring at most one rank reduction attempt per iteration, whereas P²GDR can require up to r attempts.

Observe that the $\frac{1}{2}$ factor in the sufficient decrease condition of the following result does not appear in the Armijo condition given in [7, Corollary 3.2].

COROLLARY 4.2. *Every $\tilde{X} \in \text{RFD}(X; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c)$ satisfies the sufficient decrease condition*

$$(10) \quad f(\tilde{X}) \leq f(X) - \frac{1}{2} c \alpha s(X; f, \mathbb{R}_{\leq r}^{m \times n})^2$$

for some $\alpha \in [\min\{\underline{\alpha}, 2\beta \frac{1-c}{\text{Lip}_{\mathcal{B}}(\nabla f)}\}, \bar{\alpha}]$, where \mathcal{B} is any closed ball as in Proposition 4.1. Moreover, the number of iterations in the while loop is at most

$$(11) \quad \max \left\{ 0, \left\lceil \ln \left(\frac{2(1-c)}{\alpha_0 \text{Lip}_{\mathcal{B}}(\nabla f)} \right) / \ln \beta \right\rceil \right\},$$

where α_0 is the initial step size chosen in line 1.

Proof. For all $\alpha \in (0, \infty)$,

$$f(X) + \|G\|^2 \alpha \left(-1 + \frac{\text{Lip}_{\mathcal{B}}(\nabla f)}{2} \alpha \right) \leq f(X) - c\|G\|^2 \alpha \quad \text{iff} \quad \alpha \leq 2 \frac{1-c}{\text{Lip}_{\mathcal{B}}(\nabla f)}.$$

Since the left-hand side of the first inequality is an upper bound on $f(X + \alpha G)$ for all $\alpha \in (0, \bar{\alpha}]$, the condition for the while loop to stop is necessarily satisfied if $\alpha \in (0, \min\{\bar{\alpha}, 2 \frac{1-c}{\text{Lip}_{\mathcal{B}}(\nabla f)}\}]$. Therefore, either the initial step size chosen in $[\underline{\alpha}, \bar{\alpha}]$ satisfies that condition or the while loop ends with α such that $\frac{\alpha}{\beta} > 2 \frac{1-c}{\text{Lip}_{\mathcal{B}}(\nabla f)}$. The sufficient decrease condition (10) then follows from the second inequality of (5). The upper bound (11) follows from the same observation. \square

5. The proposed algorithm. We now introduce a first-order optimization algorithm on $\mathbb{R}_{\leq r}^{m \times n}$ called RFDR. The iteration map of this algorithm, called the RFDR map, is defined as Algorithm 2. Given $X \in \mathbb{R}_{\leq r}^{m \times n}$ as input, the RFDR map proceeds as follows: (i) it applies the RFD map (Algorithm 1) to X , thereby producing a point \tilde{X} , (ii) if $\sigma_r(X)$ is positive but smaller than some threshold $\Delta \in (0, \infty)$, it applies the RFD map to a projection \hat{X} of X onto $\mathbb{R}_{r-1}^{m \times n}$, then producing a point \tilde{X}^R , and (iii) it outputs a point among \tilde{X} and \tilde{X}^R that maximally decreases f .

Algorithm 2. RFDR map.

Require: $(f, r, \underline{\alpha}, \bar{\alpha}, \beta, c, \Delta)$ where $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is differentiable with ∇f locally Lipschitz continuous, $r < \min\{m, n\}$ is a positive integer, $0 < \underline{\alpha} \leq \bar{\alpha} < \infty$, $\beta, c \in (0, 1)$, and $\Delta \in (0, \infty)$.

Input: $X \in \mathbb{R}_{\leq r}^{m \times n}$ such that $s(X; f, \mathbb{R}_{\leq r}^{m \times n}) > 0$.

Output: a point in $\text{RFDR}(X; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c, \Delta)$.

- 1: Choose $\tilde{X} \in \text{RFD}(X; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c)$;
 - 2: **if** $\sigma_r(X) \in (0, \Delta]$ **then**
 - 3: Choose $\hat{X} \in P_{\mathbb{R}_{r-1}^{m \times n}}(X)$;
 - 4: Choose $\tilde{X}^R \in \text{RFD}(\hat{X}; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c)$;
 - 5: Return $Y \in \text{argmin}_{\{\tilde{X}, \tilde{X}^R\}} f$.
 - 6: **else**
 - 7: Return \tilde{X} .
 - 8: **end if**
-

The RFDR algorithm is defined as Algorithm 3. It produces a sequence along which f is strictly decreasing. As announced in the second remark following Proposition 4.1, it turns out that the mere monitoring of the r th singular value of the iterates is enough to make the algorithm apocalypse-free, which we prove in the next section.

6. Convergence analysis. The purpose of this section is to prove Theorem 6.2. To this end, we use the abstract framework proposed in [9, section 1.3]. Indeed, the problem considered in this paper can be formulated as follows: find $X \in \mathbb{R}_{\leq r}^{m \times n}$ such

Algorithm 3. RFDR.

Require: $(X_0, f, r, \underline{\alpha}, \bar{\alpha}, \beta, c, \Delta)$ where $X_0 \in \mathbb{R}_{\leq r}^{m \times n}$, $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is differentiable with ∇f locally Lipschitz continuous, $r < \min\{m, n\}$ is a positive integer, $0 < \underline{\alpha} \leq \bar{\alpha} < \infty$, $\beta, c \in (0, 1)$, and $\Delta \in (0, \infty)$.

- 1: $i \leftarrow 0$;
 - 2: **while** $s(X_i; f, \mathbb{R}_{\leq r}^{m \times n}) > 0$ **do**
 - 3: Choose $X_{i+1} \in \text{RFDR}(X_i; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c, \Delta)$;
 - 4: $i \leftarrow i + 1$;
 - 5: **end while**
-

that $s(X; f, \mathbb{R}_{\leq r}^{m \times n}) = 0$. It is thus a particular instance of [9, Abstract Problem 1] where the Banach space is $\mathbb{R}^{m \times n}$, its closed subset is $\mathbb{R}_{\leq r}^{m \times n}$, and “desirable” means stationary. Moreover, RFDR (Algorithm 3) is a particular instance of [9, Algorithm Model 9], where the “stop rule” is f and the “search function” is the RFDR map (Algorithm 2).

Thus, to prove Theorem 6.2, it suffices to verify that RFDR satisfies the two assumptions of [9, Theorem 10], which we do below.

The first assumption is that the objective function f is continuous at each non-desirable point or bounded from below on $\mathbb{R}_{\leq r}^{m \times n}$. It is thus satisfied since f is continuous. The following proposition states that the second assumption is satisfied.

PROPOSITION 6.1. *For every $\underline{X} \in \mathbb{R}_{\leq r}^{m \times n}$ such that $s(\underline{X}; f, \mathbb{R}_{\leq r}^{m \times n}) > 0$, there exist $\varepsilon(\underline{X}), \delta(\underline{X}) \in (0, \infty)$ such that, for all $X \in B[\underline{X}, \varepsilon(\underline{X})] \cap \mathbb{R}_{\leq r}^{m \times n}$ and all $Y \in \text{RFDR}(X; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c, \Delta)$,*

$$(12) \quad f(Y) - f(X) \leq -\delta(\underline{X}).$$

Proof. Let $\underline{X} \in \mathbb{R}_{\leq r}^{m \times n}$ be such that $s(\underline{X}; f, \mathbb{R}_{\leq r}^{m \times n}) > 0$. This proof constructs $\varepsilon(\underline{X})$ and $\delta(\underline{X})$ based on the sufficient decrease condition (10) given in Corollary 4.2. This requires deriving local lower and upper bounds on $s(\cdot; f, \mathbb{R}_{\leq r}^{m \times n})$ around \underline{X} . It first considers the case where $\text{rank } \underline{X} = r$, in which the construction essentially relies on the continuity of $s(\cdot; f, \mathbb{R}_{\leq r}^{m \times n})$ on $B(\underline{X}, \sigma_r(\underline{X})) \cap \mathbb{R}_{\leq r}^{m \times n}$. Then, it focuses on the case where $\text{rank } \underline{X} < r$, in which the bounds on $s(\cdot; f, \mathbb{R}_{\leq r}^{m \times n})$ follow from the bounds (14) on ∇f thanks to (7). If $\text{rank } X < r$, then the second inequality of (7) readily gives a lower bound on $s(X; f, \mathbb{R}_{\leq r}^{m \times n})$. This is not the case if $\text{rank } X = r$, however. This is where the rank reduction mechanism comes into play. It considers a projection \hat{X} of X onto $\mathbb{R}_{r-1}^{m \times n}$, and the second inequality of (7) gives a lower bound on $s(\hat{X}; f, \mathbb{R}_{\leq r}^{m \times n})$. The inequality (12) is then obtained from (15) which follows from the continuity of f at \underline{X} .

Let us first consider the case where $\text{rank } \underline{X} = r$. On $B(\underline{X}, \sigma_r(\underline{X})) \cap \mathbb{R}_{\leq r}^{m \times n} = B(\underline{X}, \sigma_r(\underline{X})) \cap \mathbb{R}_r^{m \times n}$, $s(\cdot; f, \mathbb{R}_{\leq r}^{m \times n})$ coincides with the norm of the Riemannian gradient of the restriction of f to the smooth manifold $\mathbb{R}_r^{m \times n}$, which is continuous. In particular, there exists $\rho(\underline{X}) \in (0, \sigma_r(\underline{X}))$ such that $s(B[\underline{X}, \rho(\underline{X})] \cap \mathbb{R}_r^{m \times n}; f, \mathbb{R}_{\leq r}^{m \times n}) \subseteq [\frac{1}{2}s(\underline{X}; f, \mathbb{R}_{\leq r}^{m \times n}), \frac{3}{2}s(\underline{X}; f, \mathbb{R}_{\leq r}^{m \times n})]$. Let $\bar{\rho}(\underline{X}) := \rho(\underline{X}) + \frac{3}{2}\bar{\alpha}s(\underline{X}; f, \mathbb{R}_{\leq r}^{m \times n})$, $\varepsilon(\underline{X}) := \rho(\underline{X})$, and

$$\delta(\underline{X}) := \frac{1}{8}c s(\underline{X}; f, \mathbb{R}_{\leq r}^{m \times n})^2 \min \left\{ \underline{\alpha}, 2\beta \frac{1-c}{\text{Lip}_{B[\underline{X}, \bar{\rho}(\underline{X})]}(\nabla f)} \right\}.$$

Let $X \in B[\underline{X}, \varepsilon(\underline{X})] \cap \mathbb{R}_{\leq r}^{m \times n}$. Then, $B[X, \bar{\alpha}s(X; f, \mathbb{R}_{\leq r}^{m \times n})] \subseteq B[\underline{X}, \bar{\rho}(\underline{X})]$. Indeed, for all $Z \in B[X, \bar{\alpha}s(X; f, \mathbb{R}_{\leq r}^{m \times n})]$,

$$\|Z - \underline{X}\| \leq \|Z - X\| + \|X - \underline{X}\| \leq \bar{\alpha}s(X; f, \mathbb{R}_{\leq r}^{m \times n}) + \rho(\underline{X}) \leq \bar{\rho}(\underline{X}).$$

Therefore, Corollary 4.2 applies and, by (10), for all $\tilde{X} \in \text{RFD}(X; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c)$,

$$f(\tilde{X}) \leq f(X) - \frac{1}{2}cs(X; f, \mathbb{R}_{\leq r}^{m \times n})^2 \min \left\{ \underline{\alpha}, 2\beta \frac{1-c}{\text{Lip}_{B[\underline{X}, \bar{\rho}(\underline{X})]}(\nabla f)} \right\} \leq f(X) - \delta(\underline{X}),$$

and thus, for all $Y \in \text{RFDR}(X; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c, \Delta)$,

$$f(Y) \leq f(\tilde{X}) \leq f(X) - \delta(\underline{X}).$$

This completes the proof for the case where $\text{rank } \underline{X} = r$.

Let us now consider the case where $\text{rank } \underline{X} < r$. By the first inequality of (7), $\|\nabla f(\underline{X})\| \geq s(\underline{X}; f, \mathbb{R}_{\leq r}^{m \times n})$. Let $\bar{\rho}(\underline{X}) := \Delta + \frac{3}{2}\bar{\alpha}\|\nabla f(\underline{X})\|$ and

$$(13) \quad \delta(\underline{X}) := \frac{1}{24}c \frac{\|\nabla f(\underline{X})\|^2}{\min\{m, n\} - r + 1} \min \left\{ \underline{\alpha}, 2\beta \frac{1-c}{\text{Lip}_{B[\underline{X}, \bar{\rho}(\underline{X})]}(\nabla f)} \right\}.$$

Since f is continuous at \underline{X} , there exists $\rho_f(\underline{X}) \in (0, \infty)$ such that $f(B[\underline{X}, \rho_f(\underline{X})]) \subseteq [f(\underline{X}) - \delta(\underline{X}), f(\underline{X}) + \delta(\underline{X})]$. Let $\varepsilon(\underline{X}) := \frac{1}{2} \min\{\Delta, \rho_f(\underline{X}), \frac{\|\nabla f(\underline{X})\|}{2 \text{Lip}_{B[\underline{X}, \Delta]}(\nabla f)}\}$. Then, for all $Z \in B[\underline{X}, 2\varepsilon(\underline{X})]$, since

$$\begin{aligned} \|\nabla f(Z) - \nabla f(\underline{X})\| &\leq \|\nabla f(Z) - \nabla f(\underline{X})\| \\ &\leq \text{Lip}_{B[\underline{X}, \Delta]}(\nabla f) \|Z - \underline{X}\| \\ &\leq \text{Lip}_{B[\underline{X}, \Delta]}(\nabla f) 2\varepsilon(\underline{X}) \\ &\leq \frac{\|\nabla f(\underline{X})\|}{2}, \end{aligned}$$

it holds that

$$(14) \quad \frac{1}{2}\|\nabla f(\underline{X})\| \leq \|\nabla f(Z)\| \leq \frac{3}{2}\|\nabla f(\underline{X})\|.$$

Let $X \in B[\underline{X}, \varepsilon(\underline{X})] \cap \mathbb{R}_{\leq r}^{m \times n}$. Let us first consider the case where $\text{rank } X = r$. Then,

$$0 < \sigma_r(X) = \sigma_r(X) - \sigma_r(\underline{X}) \leq \|X - \underline{X}\| \leq \varepsilon(\underline{X}) \leq \Delta.$$

Thus, given X as input, the RFDR map considers $\hat{X} \in P_{\mathbb{R}_{r-1}^{m \times n}}(X) \subseteq B[\underline{X}, 2\varepsilon(\underline{X})] \cap \mathbb{R}_{r-1}^{m \times n}$ and $\tilde{X}^R \in \text{RFD}(\hat{X}; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c)$, where the inclusion holds because

$$\|\hat{X} - \underline{X}\| \leq \|\hat{X} - X\| + \|X - \underline{X}\| \leq \sigma_r(X) + \varepsilon(\underline{X}) \leq 2\varepsilon(\underline{X}).$$

As $X, \hat{X} \in B[\underline{X}, \rho_f(\underline{X})]$, we have

$$(15) \quad f(\hat{X}) \leq f(X) + 2\delta(\underline{X}).$$

Since $B[\hat{X}, \bar{\alpha}s(\hat{X}; f, \mathbb{R}_{\leq r}^{m \times n})] \subseteq B[\underline{X}, \bar{\rho}(\underline{X})]$, Corollary 4.2 applies to \tilde{X}^R with the ball $B[\underline{X}, \bar{\rho}(\underline{X})]$. The inclusion holds because, for all $Z \in B[\hat{X}, \bar{\alpha}s(\hat{X}; f, \mathbb{R}_{\leq r}^{m \times n})]$,

$$\|Z - \underline{X}\| \leq \|Z - \hat{X}\| + \|\hat{X} - \underline{X}\| \leq \bar{\alpha}s(\hat{X}; f, \mathbb{R}_{\leq r}^{m \times n}) + 2\varepsilon(\underline{X}) \leq \bar{\alpha}\|\nabla f(\hat{X})\| + \Delta \leq \bar{\rho}(\underline{X}),$$

the last inequality following from (14). Thus, for all $Y \in \text{RFDR}(X; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c, \Delta)$,

$$\begin{aligned} f(Y) &\leq f(\tilde{X}^R) \\ &\leq f(\hat{X}) - \frac{1}{2} c s(\hat{X}; f, \mathbb{R}_{\leq r}^{m \times n})^2 \min \left\{ \alpha, 2\beta \frac{1-c}{\text{Lip}_{B[\underline{X}, \bar{\rho}(\underline{X})]}(\nabla f)} \right\} \\ &\leq f(\hat{X}) - \frac{1}{2} c \frac{\|\nabla f(\hat{X})\|^2}{\min\{m, n\} - r + 1} \min \left\{ \alpha, 2\beta \frac{1-c}{\text{Lip}_{B[\underline{X}, \bar{\rho}(\underline{X})]}(\nabla f)} \right\} \\ &\leq f(\hat{X}) - 3\delta(\underline{X}) \\ &\leq f(X) - \delta(\underline{X}), \end{aligned}$$

where the second inequality follows from (10), the third from the second inequality of (7), the fourth from (14) and (13), and the fifth from (15). Let us now consider the case where $\text{rank } X < r$. Let $\tilde{X} \in \text{RFD}(X; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c)$. Since $B[X, \bar{\alpha} s(X; f, \mathbb{R}_{\leq r}^{m \times n})] \subseteq B[\underline{X}, \bar{\rho}(\underline{X})]$, Corollary 4.2 applies to \tilde{X} with $B[\underline{X}, \bar{\rho}(\underline{X})]$. The inclusion holds because, for all $Z \in B[X, \bar{\alpha} s(X; f, \mathbb{R}_{\leq r}^{m \times n})]$,

$$\|Z - \underline{X}\| \leq \|Z - X\| + \|X - \underline{X}\| \leq \bar{\alpha} s(X; f, \mathbb{R}_{\leq r}^{m \times n}) + \varepsilon(\underline{X}) < \bar{\alpha} \|\nabla f(X)\| + \Delta \leq \bar{\rho}(\underline{X}),$$

the last inequality following from (14). Thus, for all $Y \in \text{RFDR}(X; f, r, \underline{\alpha}, \bar{\alpha}, \beta, c, \Delta)$,

$$\begin{aligned} f(Y) &\leq f(\tilde{X}) \\ &\leq f(X) - \frac{1}{2} c s(X; f, \mathbb{R}_{\leq r}^{m \times n})^2 \min \left\{ \alpha, 2\beta \frac{1-c}{\text{Lip}_{B[\underline{X}, \bar{\rho}(\underline{X})]}(\nabla f)} \right\} \\ &\leq f(X) - \frac{1}{2} c \frac{\|\nabla f(X)\|^2}{\min\{m, n\} - r + 1} \min \left\{ \alpha, 2\beta \frac{1-c}{\text{Lip}_{B[\underline{X}, \bar{\rho}(\underline{X})]}(\nabla f)} \right\} \\ &\leq f(X) - 3\delta(\underline{X}) \\ &\leq f(X) - \delta(\underline{X}), \end{aligned}$$

where the second inequality follows from (10), the third from the second inequality of (7), and the fourth from (14) and (13). \square

We have thus proven the following.

THEOREM 6.2. *Consider a sequence constructed by RFDR (Algorithm 3). If this sequence is finite, then its last element is stationary, i.e., is a zero of the stationarity measure $s(\cdot; f, \mathbb{R}_{\leq r}^{m \times n})$ defined in (3). If it is infinite, then each of its accumulation points is stationary.*

COROLLARY 6.3. *Let $(X_i)_{i \in \mathbb{N}}$ be a sequence produced by RFDR (Algorithm 3). The sequence has at least one accumulation point if and only if $\liminf_{i \rightarrow \infty} \|X_i\| < \infty$. For every convergent subsequence $(X_{i_k})_{k \in \mathbb{N}}$, $\lim_{k \rightarrow \infty} s(X_{i_k}; f, \mathbb{R}_{\leq r}^{m \times n}) = 0$. If $(X_i)_{i \in \mathbb{N}}$ is bounded, which is the case notably if the sublevel set $\{X \in \mathbb{R}_{\leq r}^{m \times n} \mid f(X) \leq f(X_0)\}$ is bounded, then $\lim_{i \rightarrow \infty} s(X_i; f, \mathbb{R}_{\leq r}^{m \times n}) = 0$, and all accumulation points have the same image by f .*

Proof. The “if and only if” statement is a classic result. The two limits follow from [7, Proposition 2.12]. The final claim follows from the argument given in the proof of [9, Theorem 65]. \square

7. Practical implementation and computational cost. In this section, we compare the computational cost of the P²GDR map [7, Algorithm 3] with the one of the RFDR map (Algorithm 2). The comparison, summarized in Table 7.1, is conducted based on detailed implementations of these algorithms involving only evaluations of f and ∇f and some operations from linear algebra:

1. matrix multiplication;
2. thin QR factorization with column pivoting (see, e.g., [2, Algorithm 5.4.1]);
3. *small scale* (truncated) SVD, i.e., the smallest dimension of the matrix to decompose is at most $2r$;
4. *large scale* truncated SVD, i.e., truncated SVD that is not small scale.

Let us recall that, in this list, only the (truncated) SVD cannot be executed within a finite number of arithmetic operations.

Detailed implementations of the P²GDR and RFDR maps are provided as Algorithms 6 and 8, respectively. They both use Algorithm 4, which is a detailed implementation of the P²GD, P²GDR, RFD, and RFDR maps in the case where their input is $0_{m \times n}$, as a subroutine. They also rely on Algorithms 5 and 7, which are detailed implementations of the P²GD and RFD maps, respectively. All those algorithms work with low-rank representations of the involved matrices as much as possible. In the rest of the section, we analyze the computational cost of the five algorithms mentioned in this paragraph.

Algorithm 4 involves one evaluation of ∇f (in line 1), at most

$$(16) \quad 1 + \max \left\{ 0, \left\lceil \ln \left(\frac{2(1-c)}{\alpha \text{Lip}_{B[0_{m \times n}, \alpha s(0_{m \times n}; f, \mathbb{R}_{\leq r}^{m \times n})]}(\nabla f)} \right) / \ln \beta \right\rceil \right\}$$

evaluations of f (in the while loop), and one large scale truncated SVD (in line 1). The upper bound (16) follows from (11).

Algorithm 4. Detailed P²GD, P²GDR, RFD, and RFDR maps given zero as input.

Require: (f, r, α, β, c) where $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is differentiable with ∇f locally

Lipschitz continuous, $r < \min\{m, n\}$ is a positive integer, $\alpha \in (0, \infty)$, $\beta, c \in (0, 1)$, and $s(0_{m \times n}; f, \mathbb{R}_{\leq r}^{m \times n}) > 0$.

Output: $(\tilde{U}, \tilde{\Sigma}, \tilde{V})$ where $\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \mathbb{R}_{\tilde{r}}^{m \times n}$ is an SVD, $\tilde{r} \in \{1, \dots, r\}$, and

$\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \text{P}^2\text{GD}(0_{m \times n}; f, r, \alpha, \alpha, \beta, c) = \text{RFD}(0_{m \times n}; f, r, \alpha, \alpha, \beta, c) =$

$\text{P}^2\text{GDR}(0_{m \times n}; f, r, \alpha, \alpha, \beta, c, \Delta) = \text{RFDR}(0_{m \times n}; f, r, \alpha, \alpha, \beta, c, \Delta)$.

1: Compute a truncated SVD $\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \mathbb{R}_{\leq r}^{m \times n}$ of $-\nabla f(0_{m \times n})$;

2: $s \leftarrow \|\tilde{\Sigma}\|^2$;

3: **while** $f(\alpha\tilde{U}\tilde{\Sigma}\tilde{V}^\top) > f(0_{m \times n}) - c\alpha s$ **do**

4: $\alpha \leftarrow \alpha\beta$;

5: **end while**

6: $\tilde{\Sigma} \leftarrow \alpha\tilde{\Sigma}$;

Algorithm 5 involves one evaluation of ∇f (in line 1), at most

$$(17) \quad 1 + \max \left\{ 0, \left\lceil \ln \left(\frac{1-c}{\alpha \kappa_{B[X, (1+\frac{1}{\sqrt{2}})\alpha s(X; f, \mathbb{R}_{\leq r}^{m \times n})]}(X; f, \alpha)} \right) / \ln \beta \right\rceil \right\}$$

Algorithm 5. Detailed P²GD map (based on [11, Algorithms 2 and 3]).

Require: (f, r, α, β, c) where $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is differentiable with ∇f locally Lipschitz continuous, $r < \min\{m, n\}$ is a positive integer, $\alpha \in (0, \infty)$, and $\beta, c \in (0, 1)$.

Input: (U, Σ, V) where $U\Sigma V^\top \in \mathbb{R}_r^{m \times n}$ is an SVD, $r \in \{1, \dots, r\}$, and $s(U\Sigma V^\top; f, \mathbb{R}_{\leq r}^{m \times n}) > 0$.

Output: $(\tilde{U}, \tilde{\Sigma}, \tilde{V})$ where $\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \mathbb{R}_{\tilde{r}}^{m \times n}$ is an SVD, $\tilde{r} \in \{1, \dots, r\}$, and $\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \text{P}^2\text{GD}(U\Sigma V^\top; f, r, \alpha, \alpha, \beta, c)$.

- 1: $\bar{G} \leftarrow -\nabla f(U\Sigma V^\top)$; $\hat{G}_1 \leftarrow U^\top \bar{G}$; $\hat{G}_2 \leftarrow \bar{G}V$; $\hat{G} \leftarrow \hat{G}_1V$; $\hat{G}_1 \leftarrow U\hat{G}$; $\hat{G}_2 \leftarrow \hat{G}V^\top$;
 - 2: Compute QR factorizations with column pivoting $\hat{G}_2 - \hat{G}_1 = U_\perp R_1$ and $\hat{G}_1^\top - \hat{G}_2^\top = V_\perp R_2$ where $U_\perp \in \text{St}(r_1, m)$, $V_\perp \in \text{St}(r_2, n)$, $R_1 \in \mathbb{R}_*^{r_1 \times r}$, and $R_2 \in \mathbb{R}_*^{r_2 \times r}$;
 - 3: $s \leftarrow \|\hat{G}\|^2 + \|R_1\|^2 + \|R_2\|^2$;
 - 4: **if** $r = r$ **then**
 - 5: Compute a truncated SVD $\hat{U}\hat{\Sigma}\hat{V}^\top \in \mathbb{R}_{\leq r}^{r+r_1 \times r+r_2}$ of $\begin{bmatrix} \Sigma + \alpha\hat{G} & \alpha R_2^\top \\ \alpha R_1 & 0_{r_1 \times r_2} \end{bmatrix}$;
 - 6: **while** $f([U U_\perp]\hat{U}\hat{\Sigma}\hat{V}^\top[V V_\perp]^\top) > f(U\Sigma V^\top) - c\alpha s$ **do**
 - 7: $\alpha \leftarrow \alpha\beta$;
 - 8: Compute a truncated SVD $\hat{U}\hat{\Sigma}\hat{V}^\top \in \mathbb{R}_{\leq r}^{r+r_1 \times r+r_2}$ of $\begin{bmatrix} \Sigma + \alpha\hat{G} & \alpha R_2^\top \\ \alpha R_1 & 0_{r_1 \times r_2} \end{bmatrix}$;
 - 9: **end while**
 - 10: $\tilde{\Sigma} \leftarrow \hat{\Sigma}$; $\tilde{U} \leftarrow [U U_\perp]\hat{U}$; $\tilde{V} \leftarrow [V V_\perp]\hat{V}$;
 - 11: **else**
 - 12: $\tilde{G} \leftarrow \bar{G} - U\hat{G}_1 + (\hat{G}_1 - \hat{G}_2)V^\top$;
 - 13: **if** $\tilde{G} = 0_{m \times n}$ **then**
 - 14: Repeat lines 5 to 10;
 - 15: **else**
 - 16: Compute a truncated SVD $\underline{U}_\perp \underline{\Sigma} \underline{V}_\perp^\top \in \mathbb{R}_{\leq r-r}^{m \times n}$ of $\tilde{G} \in \mathbb{R}_{\leq \min\{m, n\}-r}^{m \times n}$;
 - 17: $s \leftarrow s + \|\underline{\Sigma}\|^2$; $r_0 \leftarrow \text{rank } \underline{\Sigma}$;
 - 18: Compute QR factorizations with column pivoting $[U_\perp \underline{U}_\perp] = [U_\perp \bar{U}_\perp] \begin{bmatrix} I_{r_1} & R_{1,1} \\ 0_{r_3 \times r_1} & R_{1,2} \end{bmatrix}$ and $[V_\perp \underline{V}_\perp] = [V_\perp \bar{V}_\perp] \begin{bmatrix} I_{r_2} & R_{2,1} \\ 0_{r_4 \times r_2} & R_{2,2} \end{bmatrix}$ where $\bar{U}_\perp \in \text{St}(r_3, m)$, $\bar{V}_\perp \in \text{St}(r_4, n)$, $R_{1,1} \in \mathbb{R}^{r_1 \times r_0}$, $R_{1,2} \in \mathbb{R}_*^{r_3 \times r_0}$, $R_{2,1} \in \mathbb{R}^{r_2 \times r_0}$, and $R_{2,2} \in \mathbb{R}_*^{r_4 \times r_0}$;
 - 19: Compute a truncated SVD $\hat{U}\hat{\Sigma}\hat{V}^\top \in \mathbb{R}_{\leq r}^{r+r_1+r_3 \times r+r_2+r_4}$ of $\begin{bmatrix} \Sigma + \alpha\hat{G} & \alpha R_2^\top & 0_{r \times r_4} \\ \alpha R_1 & \alpha R_{1,1} \underline{\Sigma} R_{2,1}^\top & \alpha R_{1,1} \underline{\Sigma} R_{2,2}^\top \\ 0_{r_3 \times r} & \alpha R_{1,2} \underline{\Sigma} R_{2,1}^\top & \alpha R_{1,2} \underline{\Sigma} R_{2,2}^\top \end{bmatrix}$;
 - 20: **while** $f([U U_\perp \bar{U}_\perp]\hat{U}\hat{\Sigma}\hat{V}^\top[V V_\perp \bar{V}_\perp]^\top) > f(U\Sigma V^\top) - c\alpha s$ **do**
 - 21: $\alpha \leftarrow \alpha\beta$;
 - 22: Compute a truncated SVD $\hat{U}\hat{\Sigma}\hat{V}^\top \in \mathbb{R}_{\leq r}^{r+r_1+r_3 \times r+r_2+r_4}$ of $\begin{bmatrix} \Sigma + \alpha\hat{G} & \alpha R_2^\top & 0_{r \times r_4} \\ \alpha R_1 & \alpha R_{1,1} \underline{\Sigma} R_{2,1}^\top & \alpha R_{1,1} \underline{\Sigma} R_{2,2}^\top \\ 0_{r_3 \times r} & \alpha R_{1,2} \underline{\Sigma} R_{2,1}^\top & \alpha R_{1,2} \underline{\Sigma} R_{2,2}^\top \end{bmatrix}$;
 - 23: **end while**
 - 24: $\tilde{\Sigma} \leftarrow \hat{\Sigma}$; $\tilde{U} \leftarrow [U U_\perp \bar{U}_\perp]\hat{U}$; $\tilde{V} \leftarrow [V V_\perp \bar{V}_\perp]\hat{V}$;
 - 25: **end if**
 - 26: **end if**
-

evaluations of f and small scale truncated SVDs (in the while loop that is executed), several matrix multiplications, at most four QR factorizations with column pivoting (in lines 2 and 18), and at most one large scale truncated SVD (in line 16). The upper bound (17) follows from [7, Corollary 3.2]. Let us point out that if $\hat{G}_1 = \hat{G}_2$ or $\hat{G}_2 = \hat{G}_1$ in line 2, then r_1 or r_2 is zero and empty matrices appear in the algorithm.

Based on the analysis of Algorithms 4 and 5, we analyze Algorithm 6. As mentioned in section 1, in the worst case, i.e., if the input has rank r and its largest singular value is smaller than or equal to Δ , then Algorithm 6 calls Algorithm 5 r times and Algorithm 4 once. Besides matrix multiplications, this requires $r + 1$ evaluations of ∇f , at most (16) plus r times (17) evaluations of f , at most $4r - 2$ QR factorizations with column pivoting, at most r times (17) small scale truncated SVDs, and r large scale truncated SVDs.

Algorithm 6. Detailed P²GDR map.

Require: $(f, r, \alpha, \beta, c, \Delta)$ where $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is differentiable with ∇f locally Lipschitz continuous, $r < \min\{m, n\}$ is a positive integer, $\alpha \in (0, \infty)$, $\beta, c \in (0, 1)$, and $\Delta \in (0, \infty)$.

Input: (U, Σ, V) where $U\Sigma V^\top \in \mathbb{R}_{\underline{r}}^{m \times n}$ is an SVD, $\underline{r} \in \{1, \dots, r\}$, and $s(U\Sigma V^\top; f, \mathbb{R}_{\leq \underline{r}}^{m \times n}) > 0$.

Output: $(\tilde{U}, \tilde{\Sigma}, \tilde{V})$ where $\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \mathbb{R}_{\tilde{r}}^{m \times n}$ is an SVD, $\tilde{r} \in \{1, \dots, r\}$, and $\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \text{P}^2\text{GDR}(U\Sigma V^\top; f, r, \alpha, \alpha, \beta, c, \Delta)$.

```

1:  $(\tilde{U}, \tilde{\Sigma}, \tilde{V}) \leftarrow \text{Algorithm 5}(\Sigma, U, V; f, r, \alpha, \beta, c)$ ;
2:  $r_\Delta \leftarrow |\{j \in \{1, \dots, \underline{r}\} \mid \Sigma(j, j) > \Delta\}|$ ;
3: for  $j \in \{1, \dots, \underline{r} - r_\Delta\}$  do
4:   if  $j < \underline{r}$  then
5:      $(\tilde{U}^{\text{R}}, \tilde{\Sigma}^{\text{R}}, \tilde{V}^{\text{R}}) \leftarrow \text{Algorithm 5}(U(:, 1 : \underline{r} - j), \Sigma(1 : \underline{r} - j, 1 : \underline{r} - j), V(:, 1 : \underline{r} - j); f, r, \alpha, \beta, c)$ ;
6:   else
7:      $(\tilde{U}^{\text{R}}, \tilde{\Sigma}^{\text{R}}, \tilde{V}^{\text{R}}) \leftarrow \text{Algorithm 4}(f, r, \alpha, \beta, c)$ ;
8:   end if
9:   if  $f(\tilde{U}^{\text{R}}\tilde{\Sigma}^{\text{R}}(\tilde{V}^{\text{R}})^\top) < f(\tilde{U}\tilde{\Sigma}\tilde{V}^\top)$  then
10:     $(\tilde{U}, \tilde{\Sigma}, \tilde{V}) \leftarrow (\tilde{U}^{\text{R}}, \tilde{\Sigma}^{\text{R}}, \tilde{V}^{\text{R}})$ ;
11:   end if
12: end for

```

Algorithm 7 involves one evaluation of ∇f (in line 1), at most

$$(18) \quad 1 + \max \left\{ 0, \left\lceil \ln \left(\frac{2(1-c)}{\alpha \text{Lip}_{B[X, \alpha s(X; f; \mathbb{R}_{\leq \underline{r}}^{m \times n})]}(\nabla f)} \right) / \ln \beta \right\rceil \right\}$$

evaluations of f (in the while loop that is executed), several matrix multiplications, one small scale SVD (in line 8 or 20), and at most one large scale truncated SVD (in line 15). The upper bound (18) follows from (11).

Let us mention that if one is only interested in the RFD map, then the SVD computed in line 8 or 20 is not necessary; it can be replaced, e.g., by a QR factorization with column pivoting. However, Algorithm 7 uses an SVD because it is a subroutine of Algorithm 8 which requires an SVD as input.

Algorithm 7. Detailed RFD map via SVDs (based on [11, Algorithms 2 and 4]).

Require: (f, r, α, β, c) where $f: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is differentiable with ∇f locally Lipschitz continuous, $r < \min\{m, n\}$ is a positive integer, $\alpha \in (0, \infty)$, and $\beta, c \in (0, 1)$.

Input: (U, Σ, V) where $U\Sigma V^\top \in \mathbb{R}_r^{m \times n}$ is an SVD, $r \in \{1, \dots, r\}$, and $s(U\Sigma V^\top; f, \mathbb{R}_{\leq r}^{m \times n}) > 0$.

Output: $(\tilde{U}, \tilde{\Sigma}, \tilde{V})$ where $\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \mathbb{R}_{\tilde{r}}^{m \times n}$ is an SVD, $\tilde{r} \in \{1, \dots, r\}$, and $\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \text{RFD}(U\Sigma V^\top; f, r, \alpha, \alpha, \beta, c)$.

- 1: $\tilde{G} \leftarrow -\nabla f(U\Sigma V^\top)$; $\hat{G}_1 \leftarrow U^\top \tilde{G}$; $\hat{G}_2 \leftarrow \tilde{G}V$; $s_1 \leftarrow \|\hat{G}_1\|^2$; $s_2 \leftarrow \|\hat{G}_2\|^2$;
- 2: **if** $s_1 \geq s_2$ **then**
- 3: $\hat{X}_2 \leftarrow \Sigma V^\top$;
- 4: **if** $r = r$ **then**
- 5: **while** $f(U(\hat{X}_2 + \alpha\hat{G}_1)) > f(U\Sigma V^\top) - c\alpha s_1$ **do**
- 6: $\alpha \leftarrow \alpha\beta$;
- 7: **end while**
- 8: Compute an SVD $\hat{U}\hat{\Sigma}\hat{V}^\top \in \mathbb{R}_{\hat{r}}^{r \times n}$ of $\hat{X}_2 + \alpha\hat{G}_1$;
- 9: $\tilde{U} \leftarrow U\hat{U}$;
- 10: **else**
- 11: $\tilde{G} \leftarrow \tilde{G} - U\hat{G}_1 + (U(\hat{G}_1V) - \hat{G}_2)V^\top$;
- 12: **if** $\tilde{G} = 0_{m \times n}$ **then**
- 13: Repeat lines 5 to 9;
- 14: **else**
- 15: Compute a truncated SVD $\underline{U}_\perp \underline{\Sigma} \underline{V}_\perp^\top \in \mathbb{R}_{\leq r-r}^{m \times n}$ of $\tilde{G} \in \mathbb{R}_{\leq \min\{m, n\} - r}^{m \times n}$;
- 16: $s_1 \leftarrow s_1 + \|\underline{\Sigma}\|^2$; $r_0 \leftarrow \text{rank } \underline{\Sigma}$;
- 17: **while** $f\left([U \underline{U}_\perp] \begin{bmatrix} \hat{X}_2 + \alpha\hat{G}_1 \\ \alpha \underline{\Sigma} \underline{V}_\perp^\top \end{bmatrix}\right) > f(U\Sigma V^\top) - c\alpha s_1$ **do**
- 18: $\alpha \leftarrow \alpha\beta$;
- 19: **end while**
- 20: Compute an SVD $\hat{U}\hat{\Sigma}\hat{V}^\top \in \mathbb{R}_{\hat{r}}^{r+r_0 \times n}$ of $\begin{bmatrix} \hat{X}_2 + \alpha\hat{G}_1 \\ \alpha \underline{\Sigma} \underline{V}_\perp^\top \end{bmatrix}$;
- 21: $\tilde{U} \leftarrow [U \underline{U}_\perp]\hat{U}$;
- 22: **end if**
- 23: **end if**
- 24: **else**
- 25: Repeat mutatis mutandis lines 3 to 23 with s_1 replaced by s_2 ;
- 26: **end if**

In the worst case, i.e., if the input has rank r and its smallest singular value is smaller than or equal to Δ , then, assuming that $r > 1$, Algorithm 8 calls Algorithm 7 twice. Besides matrix multiplications, this requires two evaluations of ∇f , at most two times (18) evaluations of f , two small scale SVDs, and one large scale truncated SVD.

Table 7.1 summarizes the operations required by Algorithms 6 and 8 in the case where $r > 1$, the input has rank r , and its largest singular value is smaller than or equal to Δ . We observe that Algorithm 6 requires more of each operation, except perhaps the evaluation of f , than Algorithm 8.

Algorithm 8. Detailed RFDR map.

Require: $(f, r, \alpha, \beta, c, \Delta)$ where $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ is differentiable with ∇f locally Lipschitz continuous, $r < \min\{m, n\}$ is a positive integer, $\alpha \in (0, \infty)$, $\beta, c \in (0, 1)$, and $\Delta \in (0, \infty)$.

Input: (U, Σ, V) where $U\Sigma V^\top \in \mathbb{R}_{\underline{r}}^{m \times n}$ is an SVD, $\underline{r} \in \{1, \dots, r\}$, and $s(U\Sigma V^\top; f, \mathbb{R}_{\leq \underline{r}}^{m \times n}) > 0$.

Output: $(\tilde{U}, \tilde{\Sigma}, \tilde{V})$ where $\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \mathbb{R}_{\tilde{r}}^{m \times n}$ is an SVD, $\tilde{r} \in \{1, \dots, r\}$, and $\tilde{U}\tilde{\Sigma}\tilde{V}^\top \in \text{RFDR}(U\Sigma V^\top; f, r, \alpha, \alpha, \beta, c, \Delta)$.

- 1: $(\tilde{U}, \tilde{\Sigma}, \tilde{V}) \leftarrow \text{Algorithm 7}(U, \Sigma, V; f, r, \alpha, \beta, c)$;
- 2: **if** $\underline{r} = r$ and $\Sigma(r, r) \leq \Delta$ **then**
- 3: **if** $r > 1$ **then**
- 4: $(\tilde{U}^R, \tilde{\Sigma}^R, \tilde{V}^R) \leftarrow \text{Algorithm 7}(U(:, 1:r-1), \Sigma(1:r-1, 1:r-1), V(:, 1:r-1); f, r, \alpha, \beta, c)$;
- 5: **else**
- 6: $(\tilde{U}^R, \tilde{\Sigma}^R, \tilde{V}^R) \leftarrow \text{Algorithm 4}(f, r, \alpha, \beta, c)$;
- 7: **end if**
- 8: **if** $f(\tilde{U}^R\tilde{\Sigma}^R(\tilde{V}^R)^\top) < f(\tilde{U}\tilde{\Sigma}\tilde{V}^\top)$ **then**
- 9: $(\tilde{U}, \tilde{\Sigma}, \tilde{V}) \leftarrow (\tilde{U}^R, \tilde{\Sigma}^R, \tilde{V}^R)$;
- 10: **end if**
- 11: **end if**

TABLE 7.1

Operations required by Algorithms 6 and 8 in the case where $r > 1$, the input has rank r , and its largest singular value is smaller than or equal to Δ . c. p. is column pivoting.

Operation	Algorithm 6 (P ² GDR)	Algorithm 8 (RFDR)
evaluation of f	up to (16) + $r \cdot$ (17)	up to 2 · (18)
evaluation of ∇f	$r + 1$	2
QR factorization with c. p.	up to $4r - 2$	0
small scale (truncated) SVD	$r \cdot$ (17)	2
large scale truncated SVD	r	1

8. Conclusion. We close this work with three concluding remarks.

1. As in [5] and [7], everything said in this paper remains true if f is only defined on an open subset of $\mathbb{R}^{m \times n}$ containing $\mathbb{R}_{\leq r}^{m \times n}$.
2. We stated in section 1 that RFD can follow apocalypses. It can indeed be verified that RFD follows the apocalypses described in [5, section 2.2] and [8, section 3.2]. In fact, for these two instances of (1), RFD and RFDR respectively produce the same sequences of iterates as P²GD and P²GDR. The main reason is that, for each of the two instances, if $(X_i)_{i \in \mathbb{N}}$ denotes the sequence produced by P²GD or P²GDR, then, for all $i \in \mathbb{N}$,

$$(19) \quad P_{\mathbb{R}_{\leq r}^{m \times n}(X_i)}(-\nabla f(X_i)) = P_{T_{\mathbb{R}_{\leq r}^{m \times n}(X_i)}}(-\nabla f(X_i)).$$

3. As RFD, RFDR requires the computation of at most one large scale truncated SVD per iteration. An open question is whether it is possible to design a first-order optimization algorithm on $\mathbb{R}_{\leq r}^{m \times n}$ offering the same convergence properties as RFDR without involving any large scale truncated SVD.

REFERENCES

- [1] N. GILLIS AND F. GLINEUR, *Low-rank matrix approximation with weights or missing data is NP-hard*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 1149–1165, <https://doi.org/10.1137/110820361>.
- [2] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, Johns Hopkins Stud. Math. Sci., 4th ed., Johns Hopkins University Press, Baltimore, MD, 2013.
- [3] W. HA, H. LIU, AND R. F. BARBER, *An equivalence between critical points for rank constraints versus low-rank factorizations*, SIAM J. Optim., 30 (2020), pp. 2927–2955, <https://doi.org/10.1137/18M1231675>.
- [4] J. HARRIS, *Algebraic Geometry*, Grad. Texts in Math. 133, Springer-Verlag, New York, 1992.
- [5] E. LEVIN, J. KILEEL, AND N. BOUMAL, *Finding stationary points on bounded-rank matrices: A geometric hurdle and a smooth remedy*, Math. Program., 199 (2023), pp. 831–864, <https://doi.org/10.1007/s10107-022-01851-2>.
- [6] Y. NESTEROV, *Lectures on Convex Optimization*, Springer Optim. Appl. 137, 2nd ed., Springer, Cham, 2018.
- [7] G. OLIKIER, K. A. GALLIVAN, AND P.-A. ABSIL, *An Apocalypse-Free First-Order Low-Rank Optimization Algorithm*, Technical report UCL-INMA-2022.01, <https://arxiv.org/abs/2201.03962>, 2022.
- [8] G. OLIKIER, K. A. GALLIVAN, AND P.-A. ABSIL, *Comparison of an Apocalypse-Free and an Apocalypse-Prone First-Order Low-Rank Optimization Algorithm*, Technical report UCL-INMA-2022.02, <https://arxiv.org/abs/2202.09107>, 2022.
- [9] E. POLAK, *Computational Methods in Optimization*, Math. Sci. Eng. 77, Academic Press, New York, 1971.
- [10] R. T. ROCKAFELLAR AND R. J.-B. WETS, *Variational Analysis*, Grundlehren Math. Wiss. 317, Springer-Verlag, Berlin, 1998.
- [11] R. SCHNEIDER AND A. USCHMAJEV, *Convergence results for projected line-search methods on varieties of low-rank matrices via Łojasiewicz inequality*, SIAM J. Optim., 25 (2015), pp. 622–646, <https://doi.org/10.1137/140957822>.