

École polytechnique de Louvain

Genericity of the apocalypse phenomenon in low-rank optimization

Author: **Aloïs TAVIER**

Supervisors: **Pierre-Antoine ABSIL, Guillaume OLIKIER**

Reader: **Geovani Nunes GRAPIGLIA**

Academic year 2025–2026

Master [120] in Mathematical Engineering

Abstract

Low-rank optimization aims to minimize a function over the matrices of bounded rank, a problem at the heart of data science and signal processing. A natural and inexpensive algorithm for it, known as P^2GD , is prone to a subtle failure called an *apocalypse*: it can stall at a singular point that is not a stationary point, while its own diagnostic measure falsely signals convergence. Corrective variants such as P^2GDR and $PGD-P^2GD$ repair this defect, but at a cost, and it has remained unclear whether the apocalypse is a rare accident worth tolerating or an unavoidable feature of the geometry. This thesis confronts the question on the simplest setting where the geometry can be written down explicitly, combining a systematic numerical exploration with an analysis of the continuous-time limit of P^2GD , namely its gradient flow.

Keywords: low-rank optimization, bounded-rank matrices, projected projected-gradient descent, apocalypse, stationarity, gradient flow.

Acknowledgements

I would like to sincerely thank my supervisors, Professor Pierre-Antoine Absil and Dr. Guillaume Olikier, for their guidance, support, and availability throughout this work.

I am especially grateful to Pierre-Antoine Absil, whose expertise and thoughtful suggestions greatly contributed to the development and quality of this thesis.

I am also grateful to Guillaume Olikier for introducing me to the state of the art of the field and for the many insightful discussions and suggestions that helped me better understand the subject and improve this thesis.

Finally, I would like to thank my family, friends, and loved ones for their encouragement and support.

Thank you all.

Contents

Abstract	1
1 Introduction	5
1.1 Algorithmic landscape and research question	6
1.1.1 Projected gradient descent	6
1.1.2 P²GD : cheaper iterations, no guarantee	7
1.1.3 Restored guarantees: P²GDR and PGD-P²GD	7
1.1.4 Research question	7
1.2 Contributions	8
1.3 Structure	9
2 Preliminaries	10
2.1 Notations	10
2.2 Smooth sets, singular points, and tangent spaces	11
2.3 Tangent and normal cones	12
2.4 Stationarity and the stationarity measure	13
2.5 The apocalypse	13
2.6 The P ² GD algorithm	14
3 P²GD: numerical experiments	15
3.1 Order-two: P ² GD on the double cone	15
3.1.1 The geometry of the double cone	15
3.1.2 Numerical Experiments	18
3.2 Tangent cone and projection for symmetric matrices of bounded rank	22
3.2.1 Setup and notation	22
3.2.2 Tangent and normal spaces to $S_{\tilde{r}}(n)$	23
3.2.3 Tangent cone to $S_{\leq r}(n)$	25
3.2.4 Projection onto the tangent cone	27
3.3 Order-three real symmetric matrices with bounded rank	28
3.3.1 Geometry and Parameterization of $S_{\leq 1}(3)$	28
3.3.2 Numerical Setup and Convergence Targets	28

3.3.3	Results: Persistence of the Step-Size phenomenon	29
3.3.4	Structural Observations	30
3.4	Higher-dimensional Generalizations	30
3.4.1	Experimental Protocol	31
3.4.2	Classification of Convergence Outcomes	31
3.4.3	Numerical Results and Analysis	32
4	Apocalypse Analysis: a Continuous-Time Perspective	34
4.1	The Projected Projected-Gradient Descent (P ² GD)	34
4.1.1	A Continuous-Time Model for P ² GD	34
4.1.2	The Projected Gradient Flow on the Double Cone	36
4.2	Exact Integral Equations	38
4.3	Bounds on the Trajectory	39
4.3.1	Iterative Refinement of the Bounds	39
4.3.2	First Level Bounds	40
4.3.3	Refined Lower Bound on $z(t)$	43
4.3.4	Refined Lower Bound on $\cos \theta$	43
4.3.5	Initial-Angle Criterion	44
4.4	Numerical Illustration of the Iterative Scheme	44
4.5	Global existence	49
4.5.1	Continuous-time system and basic properties	49
4.5.2	Global existence ($t_+ = \infty$)	49
4.5.3	Implication for the discrete algorithm	52
	References	56

Chapter 1

Introduction

Matrices are a central object of applied mathematics, providing a compact representation of data, linear transformations, measurements, and model parameters. In many situations, however, the high-dimensional matrix at hand reflects an underlying structure of much lower intrinsic dimension, and recovering or exploiting this structure motivates the field of low-rank optimization.

Such structure arises across a wide range of disciplines. In statistics and signal recovery, matrix completion [1], [2] and matrix sensing [3], [4] aim to reconstruct a low-rank matrix from only a limited number of measurements, reducing the amount of required data. Robust principal component analysis [5] further separates meaningful low-rank information from sparse noise or corruptions. In machine learning, multi-task learning [6] uses shared low-rank representations in order that related tasks can benefit from common information and generalize better, while weighted low-rank approximation [7], [8] accounts for data with different levels of importance or reliability. In image and video analysis, low-rank models are widely used for denoising, inpainting missing regions, and separating moving objects from static backgrounds [9]. More recently, these ideas have also been applied to deep learning: neural networks are highly effective but often require large amounts of memory and computation, motivating low-rank training approaches that reduce these costs while preserving accuracy through dynamic rank adaptation [10].

All of these problems share a common form. Given a differentiable function $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ with locally Lipschitz gradient, and a closed subset C of the determinantal variety

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

with integers m, n, r satisfying $r < \min(m, n)$, the goal is to solve

$$\min_{X \in C} f(X). \tag{1.1}$$

The set C may be the full determinantal variety, its symmetric counterpart $S_{\leq r}(n)$, or a closed subset thereof.

Solving (1.1) to global optimality is NP-hard in general: even the restricted problem of low-rank approximation with weights or missing entries is intractable in the worst case [7]. Algorithms therefore target a necessary condition for local optimality: every local minimizer of f on C is a stationary point, that is, a point $X \in C$ at which $-\nabla f(X)$ belongs to the regular normal cone $\hat{N}_C(X)$ [11, Thm. 6.12]. The formal notions of tangent cone, regular normal cone, and stationarity on the bounded-rank varieties are recalled in Chapter 2; throughout this thesis, the goal of an algorithm is to produce a sequence whose accumulation points are stationary in this sense.

1.1 Algorithmic landscape and research question

The methods proposed in the literature for (1.1) may be compared from several perspectives, including the per-iteration cost (typically dominated by the projection onto C), their mathematical sophistication, and the convergence guarantees they provide. Each algorithm involves a different trade-off among these criteria.

1.1.1 Projected gradient descent

The most direct approach is projected gradient descent (PGD [12, Algorithm 4.2]) equipped with an Armijo backtracking line search to select the step size. Starting from $X_i \in C$, we consider a trial step size $\alpha_i > 0$ and define the next iterate as

$$X_{i+1} = P_C(X_i - \alpha_i \nabla f(X_i)).$$

The Armijo condition requires that this step produces a sufficient decrease of the objective, for example

$$f(X_{i+1}) \leq f(X_i) - c \alpha_i \|\nabla f(X_i)\|^2,$$

for some fixed parameter $c \in (0, 1)$; in the backtracking procedure, α_i is computed by starting from an initial value and repeatedly multiplying it by a factor $\beta \in (0, 1)$ until this inequality is satisfied. On the full variety $\mathbb{R}_{\leq r}^{m \times n}$, Olikier and Waldspurger show that the sequence generated by such a PGD scheme has accumulation points that are Bouligand-stationary [12].

The price is the projection P_C . By Eckart–Young, $P_C(Y)$ on low-rank matrices requires the truncated singular value decomposition of Y , and $Y = X_i - \alpha \nabla f(X_i)$ is generically of full rank. Computing its truncated SVD represents a substantial cost at each iteration [13], [14]. This is the main drawback of this method despite its stationarity properties.

1.1.2 P²GD: cheaper iterations, no guarantee

The projected projected-gradient descent (P²GD) method, introduced by Schneider and Uschmajew as Algorithm 3 of [15] and extended to the bounded-rank variety in [16], eliminates the expensive full-rank SVD by restructuring the gradient step. The key observation is that C has, at each point X_i , a set of preferred directions: the directions along which one can move from X_i while remaining in C to first order. These directions form the tangent cone $T_C(X_i)$ (formally defined in Chapter 2). Instead of moving freely in $\mathbb{R}^{m \times n}$ and then projecting back onto C , P²GD first restricts the gradient to these locally feasible directions:

$$g \in P_{T_C(X_i)}(-\nabla f(X_i)), \quad X_{i+1} \in P_C(X_i + \alpha g).$$

The tangent-cone projection is efficient because $T_C(X_i)$ has an explicit structure for bounded-rank varieties (Chapter 2). Moreover, since g belongs to the tangent space of the smooth piece of C near X_i , the matrix $X_i + \alpha g$ has rank at most $2r$ [16, Sec. 6.2.2]. In contrast, $X_i - \alpha \nabla f(X_i)$ is generically full rank. Projecting a rank- $2r$ matrix onto C costs only $\mathcal{O}(r^3)$ operations via a QR factorization [13, Sec. 9.2.4] and [14, Sec. 2.3], which is substantial savings when $r \ll \min(m, n)$.

The price is the loss of convergence guarantees. Levin, Kileel, and Boumal exhibit *apocalypses* [17, Def. 2.7]: sequences whose stationarity measure tends to zero, yet whose limit is not Bouligand-stationary. An apocalypse is a false diagnostic: the algorithm declares success at a point that admits a genuine descent direction. Apocalyptic points are necessarily singular points of C , where the tangent cone enlarges.

1.1.3 Restored guarantees: P²GDR and PGD-P²GD

Two more recent algorithms restore convergence guarantees while retaining the rank- $2r$ projection of P²GD. P²GDR [16], [18] augments P²GD with an explicit rank-reduction step that re-examines lower-rank candidates when an apocalypse is suspected. PGD-P²GD [18] interleaves P²GD and PGD, using the convergence of PGD as a fallback near the singular strata. Both accumulate at Bouligand-stationary points, at the cost of a more involved algorithmic structure and analysis.

1.1.4 Research question

The corrective schemes of Section 1.1.3 are necessary in the worst case (an apocalypse *can* occur) but it is not clear whether the apocalypse is a generic obstruction, a marginal event, or a structurally stable phenomenon. If it is marginal, P²GD is enough in practice and the corrections are an over-engineered insurance; if it is structural, the corrections are indispensable.

This thesis investigates the question on the simplest non-trivial setting in which the geometry is fully explicit: the variety $S_{\leq r}(n)$ of symmetric matrices of bounded rank (see 2.1, and in particular its order-two instance $S_{\leq 1}(2)$, which is isometric to a double cone in $\mathbb{R}^3$). The choice is motivated by two facts: the tangent and normal spaces of $S_{\leq r}(n)$ admit closed-form descriptions, derived in Chapter 3 by adaptation of the rectangular analysis of [18]; and Olikier has shown, for a specific quadratic objective on $S_{\leq 1}(2)$ and a single small-step regime, that there exist two initial iterates close to one another for which P²GD generates an apocalypse [19]. Olikier also introduces a continuous-time version of P²GD on this setting, obtains a decoupled system in cylindrical coordinates, and derives partial bounds on the trajectory; the global existence of the associated gradient flow and the characterization of the basin of attraction of the apocalypse are, however, left open.

This thesis extends the work of Olikier and provides what was left open. On the theoretical side, we prove global existence of the gradient flow for all initial conditions, and derive iteratively refined upper and lower bounds that force the trajectory toward the apex for an explicit range of initial angles; we then show numerically that the discrete P²GD iterates lie within these envelopes. On the experimental side, we run P²GD from 14,000 initial conditions across a range of step sizes on $S_{\leq 1}(2)$, and extend the study to $S_{\leq 1}(3)$ and to $S_{\leq r}(n)$ for larger (n, r) . The results show that the apocalypse occupies a substantial basin of attraction in the small-step regime and persists in higher dimensions, indicating that the phenomenon is structural rather than marginal. The precise contributions are listed in Section 1.2.

1.2 Contributions

The main contributions of this thesis are the following.

- **An explicit tangent cone and projection for the symmetric bounded-rank variety.** Adapting the variational analysis of the rectangular variety [18], we obtain a closed-form description of the tangent cone to $S_{\leq r}(n)$ (Theorem 3.2) and of the orthogonal projection onto it (Proposition 3.4). These formulas make P²GD implementable on the symmetric variety in arbitrary dimension.
- **A quantitative map of the apocalypse basin of attraction.** On the double cone $S_{\leq 1}(2)$, we determine the set of initial iterates driven to an apocalypse as a function of the step size (Section 3.1). The basin grows as the step size decreases, becoming dominant in the small-step regime rather

than marginal, and the phenomenon persists on $S_{\leq 1}(3)$ (Section 3.3) and on higher-dimensional symmetric varieties (Section 3.4).

- **A continuous-time model of the small-step regime, and bounds on its trajectories.** As the step size tends to zero, P²GD is the forward-Euler discretization of the projected gradient flow $\dot{x} = P_{T_C(x)}(-\nabla f(x))$ (Section 4.1). On the double cone, this flow decouples in cylindrical coordinates into exact scalar equations for the height $z(t)$ and the angle $\theta(t)$ (Section 4.2). We derive successive upper and lower bounds on these quantities (Section 4.3): the z -coordinate is bounded from above by $z(0)e^{-t}$, which drives it toward the apex, and bounded from below by a strictly positive bound, which keeps it away from the apex in finite time. From these bounds we extract an explicit criterion on the initial angle (Section 4.3.5), and we verify numerically that the continuous trajectories lie within these envelopes, and since the discrete iterates converge to the flow as $\alpha \rightarrow 0$, this confirms the small-step apocalyptic behaviour (Section 4.4).
- **A global existence result on the continuous-time system, confirming the apocalyptic trapping.** We prove that the flow is defined for all $t \geq 0$ and every initial condition (Section 4.5), settling a question left open in [19]. The result yields the precise picture behind the experiments: the z -coordinate satisfies $z(t) > 0$ for all finite t while $z(t) \rightarrow 0$ as $t \rightarrow \infty$, so the trajectory approaches the rank-zero apex asymptotically—never reaching it in finite time, and never escaping toward the rank-one minimizer. Because this trapping holds in the step-size limit, it explains why the apocalypse dominates P²GD at small step sizes (Section 4.5.3).

1.3 Structure

This thesis is organized as follows. Chapter 2 fixes the geometric and variational background used throughout: embedded submanifolds and singular points, the tangent and normal cones, the stationarity measure, the apocalypse, and the P²GD algorithm. Chapter 3 is the numerical study: it first derives the tangent cone and the projection for $S_{\leq r}(n)$, then maps the apocalypse basin of attraction on the double cone $S_{\leq 1}(2)$, on $S_{\leq 1}(3)$, and on symmetric varieties of larger size. Chapter 4 develops the continuous-time analysis: the projected gradient flow and its exact integral equations, the iteratively refined bounds and the resulting initial-angle criterion, and finally the proof of global existence together with its implication for the discrete algorithm.

Chapter 2

Preliminaries

This chapter introduces the objects needed to define and analyze low-rank optimization problems. The thread is short: on a smooth set, classical calculus applies and a stationary point is one where the (constrained) gradient vanishes; on a set with singular points, such as the bounded-rank matrices, the tangent space must be replaced by a tangent cone, and the whole difficulty of this thesis lives at those singular points.

2.1 Notations

Throughout, m, n, r are positive integers with $r < \min(m, n)$, and E is a finite-dimensional Euclidean space with inner product $\langle \cdot, \cdot \rangle$ and norm $\| \cdot \|$. We write $\mathbb{R}^{m \times n}$ for the space of real $m \times n$ matrices, equipped with the Frobenius inner product $\langle Y, Z \rangle_F := \text{tr}(Y^\top Z)$ and norm $\| \cdot \|_F$, and $S(n) := \{Y \in \mathbb{R}^{n \times n} \mid Y^\top = Y\}$ for the symmetric matrices. Singular values are ordered $\sigma_1(X) \geq \dots \geq \sigma_{\min(m, n)}(X) \geq 0$.

The sets of main interest are the bounded-rank matrices

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

and the sets of fixed-rank matrices are

$$\mathbb{R}_r^{m \times n} := \{X \mid \text{rank}(X) = r\}, \quad S_r(n) := \{X \in S(n) \mid \text{rank}(X) = r\},$$

where the rank equals exactly r . For a closed set $C \subseteq E$ and $x \in E$, the metric projection $P_C(x) := \{y \in C \mid \|x - y\| = d(x, C)\}$ is nonempty but may contain several points. Functions $f : E \rightarrow \mathbb{R}$ are differentiable with locally Lipschitz gradient ∇f .

The sign function is

$$\operatorname{sgn} : \mathbb{R} \rightarrow \mathbb{R} : t \rightarrow \begin{cases} 1 & \text{if } t > 0 \\ -1 & \text{if } t < 0 \\ \{-1, 1\} & \text{if } t = 0 \end{cases} ;$$

it is set-valued at the origin, reflecting the multivaluedness of the projection onto the double cone encountered in Chapter 3.

We write $\mathbb{R}_+^* := (0, \infty)$ for the set of positive real numbers.

2.2 Smooth sets, singular points, and tangent spaces

A set is smooth at a point if, near that point, it looks like a flat copy of $\mathbb{R}^d$. The precise notion is the following.

Definition 2.1 (Embedded submanifold, [20, Def. 3.10]). A non-empty subset $M \subseteq E$ is a smooth embedded submanifold of E of dimension $d < \dim E$ if, for each $x \in M$, there exist a neighborhood U of x in E and a smooth function $h : U \rightarrow \mathbb{R}^{\dim E - d}$ such that

- (a) $h(y) = 0$ if and only if $y \in M \cap U$;
- (b) $\operatorname{rank} Dh(x) = \dim E - d$.

Such a function h is called a local defining function for M at x , and the tangent space to M at x is the d -dimensional linear subspace $T_M(x) := \ker Dh(x)$.

Near a smooth point, classical differential calculus on M is available, and this is the comfortable case. The sets we optimize over are not smooth everywhere: a point $x \in C$ is called singular when C is not a smooth submanifold in any neighborhood of x , and no tangent space exists in the sense above. For the bounded-rank varieties this dichotomy is explicit. The rank function is subadditive and lower semicontinuous, and its level sets of fixed rank stratify the variety: $\mathbb{R}_{\leq r}^{m \times n}$ (resp. $S_{\leq r}(n)$) is the disjoint union of its fixed-rank strata,

$$\mathbb{R}_{\leq r}^{m \times n} = \bigsqcup_{k=0}^r \mathbb{R}_k^{m \times n}, \quad S_{\leq r}(n) = \bigsqcup_{k=0}^r S_k(n),$$

each stratum $\mathbb{R}_k^{m \times n}$ (resp. $S_k(n)$) being a smooth embedded submanifold [21, Sec. 4] (see also [22, Chapter II]). A matrix is a regular point of $\mathbb{R}_{\leq r}^{m \times n}$ if and only if it lies in the top stratum, i.e. its rank equals r ; a matrix of rank $k < r$ is therefore

a singular point of $\mathbb{R}_{\leq r}^{m \times n}$, even though it is a smooth point of the stratum $\mathbb{R}_k^{m \times n}$ to which it belongs. This distinction is what makes the rank-deficient traps of Chapter 3 meaningful.

A visual example is the double cone $\mathcal{C} = \{x \in \mathbb{R}^3 : x_3^2 = x_1^2 + x_2^2\}$: it is a smooth surface everywhere except at the apex 0, its only singular point. At such points, the tangent space must be replaced by a more general notion, introduced next.

2.3 Tangent and normal cones

The generalization keeps the defining idea of a tangent space — directions along which one can move while staying, to first order, in C — and drops the requirement of linearity.

Definition 2.2 (Tangent cone, [11, Def. 6.1]). Let $C \subseteq E$ be closed and $x \in C$. The tangent cone $T_C(x)$ is the set of $v \in E$ for which there exist $x_k \rightarrow x$ in C and $t_k \rightarrow 0^+$ with

$$\frac{x_k - x}{t_k} \longrightarrow v.$$

In words, $v \in T_C(x)$ is a valid direction that can be followed from x toward nearby points of C . Properties of this cone are

- $T_C(x)$ is always a closed cone (stable under nonnegative scaling);
- at a smooth point, it coincides with the tangent space $T_M(x)$ of Section 2.2 [11, Ex. 6.8], so nothing is lost;
- at a singular point it may be nonconvex and strictly larger than the tangent space of any nearby smooth point.

Definition 2.3 (Regular normal cone, [11, Def. 6.3 and Theorem 6.28]). The regular (Fréchet) normal cone to C at x is the polar of the tangent cone,

$$\hat{N}_C(x) := \{w \in E \mid \langle w, v \rangle \leq 0 \text{ for all } v \in T_C(x)\}.$$

At a smooth point x , $T_C(x)$ is a linear subspace and $\hat{N}_C(x)$ is exactly its orthogonal complement [11, Ex. 6.8], the normal space; in that case the hat is customarily dropped and we write $N_C(x) := T_C(x)^\perp$. The orthogonal decomposition $E = T_C(x) \oplus N_C(x)$ then yields the identity used repeatedly in the sequel to compute the projection onto the tangent cone from the projection onto the normal space:

$$P_{T_C(x)}(v) = v - P_{N_C(x)}(v). \tag{2.1}$$

2.4 Stationarity and the stationarity measure

A direction $v \in T_C(x)$ is a descent direction if $\langle \nabla f(x), v \rangle < 0$, and a point $x \in C$ is stationary for f on C when it admits no descent direction in $T_C(x)$. This is equivalent to either of

$$-\nabla f(x) \in \hat{N}_C(x) \quad \text{or} \quad \left\| P_{T_C(x)}(-\nabla f(x)) \right\| = 0$$

[17, Prop. 2.5]. Every local minimizer of f on C is stationary [11, Thm. 6.12]. The last quantity above measures how far x is from stationarity and is the criterion used to monitor every algorithm in this thesis. Although $P_{T_C(x)}(\cdot)$ may be set-valued, all its elements share the same norm, so the stationarity measure is well defined.

Definition 2.4 (Stationarity measure, [17]). The stationarity measure of f on C is

$$s(x; f, C) := \left\| P_{T_C(x)}(-\nabla f(x)) \right\|.$$

An iterate x_k satisfying $s(x_k; f, C) \leq \varepsilon$ for a prescribed tolerance $\varepsilon > 0$ is called ε -stationary.

2.5 The apocalypse

On a smooth set, $s(x_k) \rightarrow 0$ guarantees that every limit point of (x_k) is stationary [18, Cor. 2.11]. On a set with singular points this can fail, and the failure is the phenomenon at the center of this thesis.

Definition 2.5 (Apocalypse, [17, Def. 2.7]). A point $x \in C$ is *apocalyptic* if there exist a sequence $(x_k)_{k \in \mathbb{N}} \subseteq C$ converging to x , a continuously differentiable function f , and a sequence $\varepsilon_k > 0$ with $\varepsilon_k \rightarrow 0$ such that each x_k is ε_k -stationary while x is not stationary. The triplet $(x, (x_k)_{k \in \mathbb{N}}, f)$ is called an *apocalypse*.

A reason for this phenomenon is the augmentation of the tangent cone noted after Definition 2.2: as smooth iterates $x_k \rightarrow x$ approach a singular limit, $T_C(x)$ acquires directions absent from each $T_C(x_k)$. If $-\nabla f(x)$ has a nonzero component along one of them, then x admits a descent direction and is not stationary, even though $s(x_k) \rightarrow 0$ has declared success. Apocalyptic points are therefore necessarily singular [17, Rem. 2.9]; for the bounded-rank varieties studied here, this links the phenomenon to the vanishing of the smallest relevant singular value, $\sigma_r(x_k) \rightarrow 0$.

The remainder of this thesis investigates this phenomenon for the P²GD algorithm: Chapter 3 measures how often it occurs as a function of the step size, and Chapter 4 explains, through a continuous-time analysis, why it dominates in the small-step-size regime.

2.6 The P²GD algorithm

The algorithm whose behavior this thesis analyzes is the projected projected-gradient descent (P²GD), introduced by Schneider and Uschmajew as Algorithm 3 of [15] and extended to stratified sets by Olikier, Gallivan and Absil [16]; Section 1.1 gives the context and the comparison with PGD.

Given the current iterate $x \in C$, P²GD first projects the negative gradient onto the tangent cone, $g \in P_{T_C(x)}(-\nabla f(x))$, then performs an Armijo backtracking line search along g , accepting a step size $\alpha \in \mathbb{R}_+^*$ only when it yields sufficient decrease relative to the stationarity measure $s(x; f, C)$ of Definition 2.4. This is used in the P²GD map (Algorithm 1). The full method iterates it until the stationarity measure is sufficiently small (Algorithm 2).

Algorithm 1 P²GD map

Require: $(E, C, f, \underline{\alpha}, \bar{\alpha}, \beta, c)$ where E is a Euclidean vector space, C a nonempty closed subset of E , $f : E \rightarrow \mathbb{R}$ is differentiable with ∇f locally Lipschitz continuous, $0 < \underline{\alpha} \leq \bar{\alpha} < \infty$, and $\beta, c \in (0, 1)$.

Input: $x \in C$ such that $s(x; f, C) > 0$.

Output: $y \in \text{P}^2\text{GD}(x; E, C, f, \underline{\alpha}, \bar{\alpha}, \beta, c)$.

- 1: Choose $g \in P_{T_C(x)}(-\nabla f(x))$, $\alpha \in [\underline{\alpha}, \bar{\alpha}]$, and $y \in P_C(x + \alpha g)$;
 - 2: **while** $f(y) > f(x) - c \alpha s(x; f, C)^2$ **do**
 - 3: $\alpha \leftarrow \alpha \beta$;
 - 4: Choose $y \in P_C(x + \alpha g)$;
 - 5: **end while**
 - 6: **return** y ;
-

Algorithm 2 P²GD

Require: $(E, C, f, \underline{\alpha}, \bar{\alpha}, \beta, c)$ where E is a Euclidean vector space, C a nonempty closed subset of E , $f : E \rightarrow \mathbb{R}$ is differentiable with ∇f locally Lipschitz continuous, $0 < \underline{\alpha} \leq \bar{\alpha} < \infty$, and $\beta, c \in (0, 1)$.

Input: $x_0 \in C$.

Output: a sequence in C .

- 1: $i \leftarrow 0$;
 - 2: **while** $s(x_i; f, C) > 0$ **do**
 - 3: Choose $x_{i+1} \in \text{P}^2\text{GD}(x_i; E, C, f, \underline{\alpha}, \bar{\alpha}, \beta, c)$;
 - 4: $i \leftarrow i + 1$;
 - 5: **end while**
-

Chapter 3

P²GD: numerical experiments

This chapter presents numerical experiments analyzing the P²GD algorithm on the set of real symmetric matrices of bounded rank. The objective is to characterize the conditions under which P²GD generates an apocalypse.

This chapter begins with an examination of order-two and order-three symmetric matrices of rank at most one, as these simplest cases of interest admit explicit geometric interpretations. Higher-dimensional generalizations of symmetric matrices of bounded rank are then investigated numerically to identify persistent patterns.

3.1 Order-two: P²GD on the double cone

3.1.1 The geometry of the double cone

The numerical study is initiated with the consideration of the simplest non-trivial case of 2×2 real symmetric matrices of rank at most one, i.e. $S_{\leq 1}(2)$ that can be formulated as,

$$S_{\leq 1}(2) = \{X \mid X^T = X \text{ and } \det(X) = 0\}.$$

This set can be partitioned into a smooth manifold $S_1(2)$ containing all rank-one matrices, and a singular point $S_0(2) = \{0_{2 \times 2}\}$. According to [23], the smooth part $S_r(n)$ of the variety of symmetric matrices of rank at most r consists of $r + 1$ connected components. In our case (rank-one matrices), this implies that $S_1(2)$ is composed of two distinct connected components which are linked by the singularity. The set $S_{\leq 1}(2)$ can be identified with a double cone

$$C = \{x \in \mathbb{R}^3 \mid x_3^2 = x_1^2 + x_2^2\}$$

as $S_{\leq 1}(2) = \varphi(C)$ via the isomorphism

$$\varphi : \mathbb{R}^3 \rightarrow \mathbb{R}^{2 \times 2} : x \mapsto \frac{1}{\sqrt{2}} \begin{pmatrix} x_3 - x_1 & x_2 \\ x_2 & x_3 + x_1 \end{pmatrix}.$$

The double cone highlights the structure of $S_{\leq 1}(2)$ each cone (without their apex) corresponding to a component of $S_1(2)$.

For every point $x \in C$ there exist an associated matrix $A = \varphi(x)$. The eigenvalues of A are

$$\lambda_1 = 0, \quad \lambda_2 = \sqrt{2}x_3.$$

It follows that for every $x \in C$, if x_3 is strictly positive then the associated matrix A is positive-semidefinite (with $\lambda_1 = 0$ and $\lambda_2 > 0$). On the other hand, if x_3 is negative, the associated matrix A will always be negative-semidefinite (with $\lambda_1 = 0$ and $\lambda_2 < 0$). The double cone C is particularly convenient to give a visual interpretation of the definiteness of a matrix in $S_{\leq 1}(2)$, summarized in Table 3.1.




Condition	Eigenvalues	$\mathbf{x}_3$	Geometric Region
$A \succeq 0$	$(0, \sqrt{2}x_3)$	$x_3 > 0$	
$A \preceq 0$	$(\sqrt{2}x_3, 0)$	$x_3 < 0$	
$A = 0_{2 \times 2}$	$(0, 0)$	$x_3 = 0$	

Table 3.1: Classification on the double cone $S_{\leq 1}(2)$ by eigenvalues.

The set C is a locally smooth embedded submanifold of $\mathbb{R}^3$ everywhere except at the origin. Let $U = \mathbb{R}^3 \setminus \{(0, 0, 0)\}$, which is an open subset of $\mathbb{R}^3$, hence a smooth manifold in its own right. Consider the smooth map:

$$F : U \rightarrow \mathbb{R}, \quad F(x_1, x_2, x_3) = x_1^2 + x_2^2 - x_3^2.$$

The punctured double cone is exactly the level set $F^{-1}(0)$. The gradient of F is given by:

$$\nabla F(x_1, x_2, x_3) = (2x_1, 2x_2, -2x_3),$$

which never vanishes on U . Since the gradient is non-zero everywhere on U , the differential dF_x is a non-zero linear form at every point, and is therefore surjective onto $\mathbb{R}$. This implies that F is a smooth submersion on U . According to the Submersion Level Set Theorem [24, Cor. 5.13], the level set $F^{-1}(0) = C \setminus \{0\}$ is a properly embedded submanifold of U (and thus of $\mathbb{R}^3$). As Remark 2.9 in [17] establishes, apocalypses only occur at singular points (where the set is locally non-smooth). Consequently, for all experiments in section 3.1, an apocalypse will be declared if, beyond a certain iteration index, all points generated by Algorithm 2 are ϵ -stationary and remain within an ϵ -neighborhood of the singular point, namely the origin.

Projections

The P²GD algorithm 2 requires both a projection onto C and a projection onto the tangent cone of C at any point $x \in C$. These projections are computed following the formulas provided in proposition 2 and 1 of [19] respectively. The projection for all $z \in \mathbb{R}^3 \setminus \{0\}$ such that $|z_3| < \sqrt{z_1^2 + z_2^2}$ is

$$P_C(z) = \frac{z}{2} + \frac{1}{2} \begin{bmatrix} \frac{|z_3|}{\sqrt{z_1^2 + z_2^2}} z_1 \\ \frac{|z_3|}{\sqrt{z_1^2 + z_2^2}} z_2 \\ \text{sgn}(z_3) \sqrt{z_1^2 + z_2^2} \end{bmatrix}, \quad (3.1)$$

The projection for all $v \in \mathbb{R}^3$ onto the tangent space of $C \setminus \{0\}$ at x is

$$P_{T_C(x)}(v) = v - P_{N_C(x)}(v). \quad (3.2)$$

Given that $N_C(x)$ is spanned by a single vector $n(x)$ (as shown in Proposition 1 of [19]), the projection onto the normal space is

$$P_{N_C(x)}(v) = \frac{\langle n(x), v \rangle}{\|n(x)\|^2} n(x), \quad \text{with} \quad n(x) = \begin{bmatrix} x_1 \\ x_2 \\ -x_3 \end{bmatrix}$$

At the origin, the tangent cone is $T_C(0) = C$ (Proposition 1 in [19]). Therefore, the above projection onto C is the projection onto the tangent space of C at the origin.

3.1.2 Numerical Experiments

To investigate the convergence behavior of P²GD (Algorithm 2), let the objective function be the quadratic cost function

$$f : \mathbb{R}^3 \rightarrow \mathbb{R} : x \mapsto \frac{1}{2} \|x - x^*\|^2, \text{ with } x^* = \begin{bmatrix} -1 \\ 0 \\ -1 \end{bmatrix}.$$

The cost function is convex with gradient $\nabla f(x) = x - x^*$ and reaches its global minimum at $x^* \in C$, which is associated with the matrix

$$\varphi(x^*) = \begin{bmatrix} 0 & 0 \\ 0 & -\sqrt{2} \end{bmatrix}$$


The parameters of Algorithm 2 are chosen as follows :

1. Step size: $\alpha = 0.4$
2. Armijo condition: $c = 10^{-4}$
3. Backtracking factor: $b = 0.5$
4. Convergence criterion: $\epsilon = 10^{-2}$

Across all experiments in Section 3.1.2 the same objective function is kept and the parameters β, c and ϵ remain constant.

Analysis of Trajectories

The first analysis of the behavior of P²GD focuses on two representative initial points: $x_{0,\min} = (1, 0, 1)$ which is associated with the positive-semidefinite matrix $\varphi(x_{0,\min}) = \begin{bmatrix} 0 & 0 \\ 0 & \sqrt{2} \end{bmatrix}$ and $x_{0,\text{apo}} = (-1, 0, 1)$ which is associated with the positive-semidefinite matrix $\varphi(x_{0,\text{apo}}) = \begin{bmatrix} \sqrt{2} & 0 \\ 0 & 0 \end{bmatrix}$. These points are chosen because their respective trajectories remain entirely within the plane $y = 0$. Indeed, for any point x belonging to this plane, the gradient of the function at this point also belongs to the plane. From the projection formulas (Equations (3.1) and (3.2)), it follows that both projection operators preserve the plane $y = 0$. Consequently, if the iterates are initialized in this plane, all subsequent points generated by the algorithm remain in $y = 0$. Within this plane, the double cone C reduces to the union of two intersecting lines $z = \pm x$. Consequently, the local tangent space at any point $x \neq 0$ coincides with the line itself, and the projection operators simplify



(a) Convergence to the global minimum. (b) Convergence to an apocalypse.

Figure 3.1: Two distinct behaviors of P^2GD on the double cone. The red points are the P^2GD iterates $(x_i)_{i=1}^{\infty}$. The green point is the global minimum. The yellow arrows represent $x_i - \alpha_i \nabla f(x_i)$ and the green arrow represent $x_{i+1} - x_i$, where $\alpha_i = 0.4$

to orthogonal projections onto these lines. This two-dimensional reduction provides a clear framework to visualize the convergence properties.

For the initial iterate $x_{0,\min} = (1, 0, 1)$, the iterations are shown in Figure 3.1a. The negative gradient $-\nabla f(x)$ (equal to the projected gradient G in this case) remains on the cone and is on the branch containing the global minimum. As in this planar case, $s(x) = \|G\| = \|\nabla f(x)\| = \|x - x^*\|$, the stationarity measure converges to zero as the sequence approaches the global minimizer.

Conversely, for $x_{0,\text{apo}} = (-1, 0, 1)$, the iterations are represented in Figure 3.1b the algorithm becomes trapped in a branch leading toward the singularity at the origin. For any iterate x_i , the projected anti-gradient G , points exactly toward the origin, with $\|G\| = \|x - 0\|$. Hence $x + \alpha G = (1 - \alpha)x$ lies on the same branch for any $\alpha \in (0, 1)$. The iterates therefore approach 0 monotonically but never reach it. As the iterates approach $x = 0$, the anti-gradient direction $-\nabla f(x)$ becomes increasingly orthogonal to the cone's surface. Consequently, the norm of the projected gradient G (i.e. the stationarity measure) tends to zero without convergence towards the global minimum.

Analysis over a grid of initial points and the impact of the step size α

To characterize the global influence of the singularity on the convergence of P^2GD , a systematic exploration of initial conditions on the upper nappe of the double

cone is performed. The upper nappe is defined as

$$C_+ = \{x \in C \mid x_3 > 0\}.$$

The domain is discretized using a cylindrical coordinate system (r, θ, z) , where $x_1 = r \cos \theta$, $x_2 = r \sin \theta$, and $x_3 = r$. The sampling of the domain was set with $r \in [0, 10]$ with 70 equidistant levels, and for each level, 200 equally spaced values of $\theta \in [-\pi, \pi]$ were chosen. This results in 14,000 initial iterates $x_0 \in C_+$.

To study the sensitivity of the apocalypse phenomenon to the step size, this experiment is reproduced for various values of α chosen in the interval $[0.1, 0.75]$.

For each initial iterate x_0 and each step size α , the algorithm terminates at an iterate X_T when the stationarity measure $s(X_T; f, C)$ becomes smaller than ϵ or the maximum iterations is reached. An initial iterate x_0 is classified as belonging to the *apocalypse basin* if it satisfies the following two conditions:

1. The stationarity measure is sufficiently small: $s(X_T; f, C) < \epsilon$.
2. The final iterate is close to the singularity: $\|X_T\| < \epsilon$.

If these conditions are met, the trajectory is considered to have converged to an apocalypse, characterizing the point x_0 as trapped by the singular geometry at the origin. As $-\nabla f(0) = x^*$ which is on the tangent cone at the origin (as $T_C(0) = C$). The norm of the projected gradient at the origin is nonzero, which is consistent with the definition of the apocalypse 2.5.

The results illustrated in Figures 3.2 and 3.3 highlight the role of the step size α on the convergence properties of the P²GD algorithm.

- **High Step Sizes** ($\alpha \in [0.6, 0.7]$): For larger values of α , the basin of attraction of the apocalypse (represented in red) exhibits a complex structure. The algorithm shows high sensitivity to initial conditions, where small perturbations in the starting point can lead to the convergence to the minimum or to the singularity.
- **Intermediate Step Sizes** ($\alpha \approx 0.55$): As α decreases, the red zones expand and merge.
- **Low Step Sizes** ($\alpha \leq 0.45$): For smaller step sizes, the basin of attraction of the apocalypse seems to dominate almost the entire feasible domain.

In conclusion, while a smaller step size usually ensures smoother convergence in standard convex optimization, on the double cone, it increases the probability of being trapped in the singularity, effectively shrinking the basin of attraction of the global minimum to a single line $L = \{x \in C \mid x_1 = x_3\}$.







Figure 3.2: Comparison of the basin of attraction of the apocalypse for various step sizes α . In red are all initial points converging to the apocalypse using P^2GD . In green is the global minimum.




Figure 3.3: Top-down orthographic projection onto the (x, y) -plane of the basin of attraction of the apocalypse for various step sizes α . The global minimum (green dot at coordinates $(-1, 0)$ near the origin) is included in the two-dimensional projection for reference, regardless of its original z -coordinate. In red are all initial points converging to the apocalypse with P²GD.

3.2 Tangent cone and projection for symmetric matrices of bounded rank

While the 2×2 case admits a direct geometric interpretation via the double cone, analyzing the P²GD algorithm in higher dimensions requires an algebraic approach. In this section, we derive an explicit description of the tangent cone to $S_{\leq r}(n)$ together with the orthogonal projection onto this cone. Our derivation adapts the variational proof of [21, Theorem 4.1], which establishes the analogous result for the rectangular variety $\mathbb{R}_{\leq r}^{m \times n}$.

3.2.1 Setup and notation

The ambient space $S(n) := \{Y \in \mathbb{R}^{n \times n} \mid Y^\top = Y\}$ of symmetric matrices is equipped with the Frobenius inner product $\langle Y, Z \rangle_F := \text{tr}(Y^\top Z)$ and the associated norm $\|\cdot\|_F$. The *Stiefel manifold* $\text{St}_{\mathbb{R}}(k, n) := \{U \in \mathbb{R}^{n \times k} \mid U^\top U = I_k\}$ is compact: it is bounded since $\|U\|_F^2 = \text{tr}(U^\top U) = k$ for all $U \in \text{St}_{\mathbb{R}}(k, n)$, and closed as the

preimage of $\{I_k\}$ under the continuous map $U \mapsto U^\top U$.

Throughout, fix $X \in S_{\leq r}(n)$ and set $\bar{r} := \text{rank}(X)$. Let $U \in \text{St}_{\mathbb{R}}(\bar{r}, n)$ and $U_\perp \in \text{St}_{\mathbb{R}}(n - \bar{r}, n)$ satisfy $\text{im}(U) = \text{im}(X)$ and $\text{im}(U_\perp) = \ker(X)$, in order that $[U \ U_\perp] \in \mathbb{R}^{n \times n}$ is orthogonal. For any $Z \in \mathbb{R}^{n \times n}$, we denote the block decomposition

$$[U \ U_\perp]^\top Z [U \ U_\perp] = \begin{bmatrix} Z_{11} & Z_{12} \\ Z_{21} & Z_{22} \end{bmatrix}, \quad (3.3)$$

with $Z_{11} \in \mathbb{R}^{\bar{r} \times \bar{r}}$, $Z_{12} \in \mathbb{R}^{\bar{r} \times (n - \bar{r})}$, and $Z_{22} \in \mathbb{R}^{(n - \bar{r}) \times (n - \bar{r})}$. In this basis, X itself reads

$$[U \ U_\perp]^\top X [U \ U_\perp] = \begin{bmatrix} \Lambda & 0 \\ 0 & 0 \end{bmatrix}, \quad (3.4)$$

where $\Lambda := U^\top X U \in S(\bar{r})$ is invertible since $\text{rank}(\Lambda) = \text{rank}(X) = \bar{r}$.

3.2.2 Tangent and normal spaces to $S_{\bar{r}}(n)$

The following proposition shows that $S_{\bar{r}}(n)$ is a smooth embedded submanifold of $S(n)$ near X , and gives explicit expressions for its tangent and normal spaces. The proof uses the regular level set theorem applied to a defining map built from the Schur complement.

Proposition 3.1. *The set $S_{\bar{r}}(n)$ is a smooth embedded submanifold of $S(n)$ near X , with tangent and normal spaces*

$$T_{S_{\bar{r}}(n)}(X) = \left\{ [U \ U_\perp] \begin{bmatrix} A & B \\ B^\top & 0 \end{bmatrix} [U \ U_\perp]^\top \mid \begin{array}{l} A \in S(\bar{r}), \\ B \in \mathbb{R}^{\bar{r} \times (n - \bar{r})} \end{array} \right\}, \quad (3.5)$$

$$N_{S_{\bar{r}}(n)}(X) = \left\{ [U \ U_\perp] \begin{bmatrix} 0 & 0 \\ 0 & C \end{bmatrix} [U \ U_\perp]^\top \mid C \in S(n - \bar{r}) \right\}. \quad (3.6)$$

The orthogonal projection onto the normal space is, for $Z \in S(n)$,

$$P_{N_{S_{\bar{r}}(n)}(X)}(Z) = U_\perp (U_\perp^\top Z U_\perp) U_\perp^\top. \quad (3.7)$$

Proof. Local defining map. The set

$$\mathcal{U} := \{Y \in S(n) \mid Y_{11} \text{ is invertible}\}$$

is open in $S(n)$, as the preimage of $\text{GL}_{\bar{r}}(\mathbb{R})$ under the continuous linear map $Y \mapsto U^\top Y U$, and contains X since Λ is invertible. Define

$$\Psi : \mathcal{U} \rightarrow S(n - \bar{r}), \quad \Psi(Y) := Y_{22} - Y_{21} Y_{11}^{-1} Y_{12}.$$

The map Ψ is smooth on $\mathcal{U}$ as a composition of linear projections, matrix products, and matrix inversion (smooth on $\text{GL}_{\bar{r}}(\mathbb{R})$). It takes values in $S(n - \bar{r})$: symmetry of Y gives $Y_{21} = Y_{12}^\top$ and Y_{11}^{-1} symmetric, hence $\Psi(Y)^\top = Y_{22} - Y_{21}Y_{11}^{-1}Y_{12} = \Psi(Y)$.

Zero set. For $Y \in \mathcal{U}$, the Schur complement factorization

$$\begin{bmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{22} \end{bmatrix} = \begin{bmatrix} I & 0 \\ Y_{21}Y_{11}^{-1} & I \end{bmatrix} \begin{bmatrix} Y_{11} & 0 \\ 0 & \Psi(Y) \end{bmatrix} \begin{bmatrix} I & Y_{11}^{-1}Y_{12} \\ 0 & I \end{bmatrix}$$

yields $\text{rank}(Y) = \text{rank}(Y_{11}) + \text{rank}(\Psi(Y)) = \bar{r} + \text{rank}(\Psi(Y))$, hence

$$Y \in S_{\bar{r}}(n) \cap \mathcal{U} \iff \Psi(Y) = 0. \quad (3.8)$$

Differential of Ψ . For $Y \in \mathcal{U}$ and $\delta Y \in S(n)$, the first-order expansion of the inverse

$$\begin{aligned} (Y_{11} + \delta Y_{11})^{-1} &= \left(Y_{11}(I + Y_{11}^{-1}\delta Y_{11}) \right)^{-1} \\ &= Y_{11}^{-1} - Y_{11}^{-1}\delta Y_{11}Y_{11}^{-1} + O(\|\delta Y_{11}\|^2), \end{aligned}$$

combined with the expansion of $\Psi(Y + \delta Y)$, gives, after collecting terms linear in δY ,

$$D\Psi_Y(\delta Y) = \delta Y_{22} - \delta Y_{21}Y_{11}^{-1}Y_{12} - Y_{21}Y_{11}^{-1}\delta Y_{12} + Y_{21}Y_{11}^{-1}\delta Y_{11}Y_{11}^{-1}Y_{12}. \quad (3.9)$$

At $Y = X$, the off-diagonal and $(2, 2)$ blocks vanish (cf. (3.4)), so the last three terms drop out:

$$D\Psi_X(\delta Y) = \delta Y_{22}. \quad (3.10)$$

Surjectivity. Let $Y \in \mathcal{U}$ and $C \in S(n - \bar{r})$. Define $\delta Y \in S(n)$ by $\delta Y_{11} = \delta Y_{12} = \delta Y_{21} = 0$ and $\delta Y_{22} = C$, i.e.,

$$\delta Y := [U \ U_\perp] \begin{bmatrix} 0 & 0 \\ 0 & C \end{bmatrix} [U \ U_\perp]^\top.$$

Substituting in (3.9), all terms but the first vanish, giving $D\Psi_Y(\delta Y) = C$. Hence $D\Psi_Y$ is surjective for every $Y \in \mathcal{U}$, so Ψ is a smooth submersion on $\mathcal{U}$.

Conclusion via the regular level set theorem. The map $\Psi : \mathcal{U} \rightarrow S(n - \bar{r})$ is a smooth submersion between finite-dimensional smooth manifolds, with 0 a regular value. By the regular level set theorem [24, Corollary 5.14], $\Psi^{-1}(0)$ is a smooth embedded submanifold of $\mathcal{U}$. Combined with (3.8), $S_{\bar{r}}(n) \cap \mathcal{U}$ is therefore a smooth embedded submanifold of $S(n)$ near X , and Ψ is a *local defining map* for $S_{\bar{r}}(n)$ at

X in the sense of [24, before Proposition 5.38]. By [24, Proposition 5.38] applied to Ψ , together with (3.10),

$$T_{S_{\bar{r}}(n)}(X) = \ker D\Psi_X = \{\delta Y \in S(n) \mid \delta Y_{22} = 0\},$$

which is exactly (3.5). The orthogonal complement of (3.5) in $S(n)$ for the Frobenius inner product is, by block decomposition, (3.6). Finally, for $Z \in S(n)$, the projection $P_{N_{S_{\bar{r}}(n)}(X)}(Z)$ retains the $(2, 2)$ block of Z and zeros the others, giving

$$P_{N_{S_{\bar{r}}(n)}(X)}(Z) = [U \ U_{\perp}] \begin{bmatrix} 0 & 0 \\ 0 & Z_{22} \end{bmatrix} [U \ U_{\perp}]^{\top} = U_{\perp}(U_{\perp}^{\top} Z U_{\perp}) U_{\perp}^{\top}. \quad \square$$

3.2.3 Tangent cone to $S_{\leq r}(n)$

The following theorem and its proof adapt [21, Theorem 4.1 and Section 5] from the rectangular variety $\mathbb{R}_{\leq r}^{m \times n}$ to its symmetric counterpart $S_{\leq r}(n)$. The adaptation amounts to setting $V = U$ and $V_{\perp} = U_{\perp}$ in their proof, which is valid since $\text{im}(X^{\top}) = \text{im}(X)$ for symmetric X .

Theorem 3.2. *With the notation above,*

$$T_{S_{\leq r}(n)}(X) = T_{S_{\bar{r}}(n)}(X) + (N_{S_{\bar{r}}(n)}(X) \cap S_{\leq r-\bar{r}}(n)) \quad (3.11)$$

$$= T_{S_{\bar{r}}(n)}(X) + S_{\leq r-\bar{r}}(n) \quad (3.12)$$

$$= \{Z \in S(n) \mid \text{rank}(U_{\perp}^{\top} Z U_{\perp}) \leq r - \bar{r}\}. \quad (3.13)$$

Proof. If $X = 0$, then $S_{\leq r}(n)$ is a closed cone in $S(n)$, so $T_{S_{\leq r}(n)}(0) = S_{\leq r}(n)$ by [21, Proposition 2.1], matching all three expressions with $\bar{r} = 0$ and $U_{\perp} = I_n$. Assume $X \neq 0$.

Equivalence of (3.11) and (3.13). For $Z \in S(n)$, the orthogonal decomposition $Z = P_{T_{S_{\bar{r}}(n)}(X)}(Z) + P_{N_{S_{\bar{r}}(n)}(X)}(Z)$ shows that Z belongs to (3.11) if and only if $P_{N_{S_{\bar{r}}(n)}(X)}(Z) \in S_{\leq r-\bar{r}}(n)$, that is, has rank at most $r - \bar{r}$. By (3.7) and the orthonormality of $U_{\perp}$, $\text{rank}(P_{N_{S_{\bar{r}}(n)}(X)}(Z)) = \text{rank}(U_{\perp}^{\top} Z U_{\perp})$, which is exactly the condition in (3.13). The inclusion (3.11) $\subseteq$ (3.12) is immediate. It remains to prove

$$(I) \quad T_{S_{\leq r}(n)}(X) \subseteq (3.13), \quad (II) \quad (3.12) \subseteq T_{S_{\leq r}(n)}(X).$$

Step (I). This argument is the symmetric counterpart of [21, Section 5.1.2]. Let $Z \in T_{S_{\leq r}(n)}(X)$. By the definition of the tangent cone [11, Definition 6.1], there exist $(X_i)_{i \in \mathbb{N}} \subset S_{\leq r}(n)$ with $X_i \rightarrow X$ and $(t_i)_{i \in \mathbb{N}} \subset (0, \infty)$ with $t_i \rightarrow 0$ such that $(X_i - X)/t_i \rightarrow Z$.

Since $\text{rank}(X_i) \leq r$, we have $\dim \ker(X_i) \geq n - r$, so we may choose $U_{i\perp} \in \text{St}_{\mathbb{R}}(n - r, n)$ with $\text{im}(U_{i\perp}) \subseteq \ker(X_i)$, hence $U_{i\perp}^\top X_i = 0$. Combined with $XU_\perp = 0$:

$$U_{i\perp}^\top \frac{X_i - X}{t_i} U_\perp = \frac{1}{t_i} (U_{i\perp}^\top X_i U_\perp - U_{i\perp}^\top X U_\perp) = 0 \quad \forall i. \quad (3.14)$$

Since $\text{St}_{\mathbb{R}}(n - r, n)$ is compact, after extracting a subsequence (still denoted by i), $U_{i\perp} \rightarrow \bar{U}_\perp \in \text{St}_{\mathbb{R}}(n - r, n)$, and the convergences $X_i \rightarrow X$ and $(X_i - X)/t_i \rightarrow Z$ are preserved. Passing to the limit in $U_{i\perp}^\top X_i = 0$ and in (3.14):

$$\bar{U}_\perp^\top X = 0, \quad \bar{U}_\perp^\top Z U_\perp = 0. \quad (3.15)$$

The first identity gives $\text{im}(\bar{U}_\perp) \subseteq \ker(X) = \text{im}(U_\perp)$, so $\bar{U}_\perp = U_\perp U_\perp^\top \bar{U}_\perp$. Set $W := U_\perp^\top \bar{U}_\perp$. Then $W^\top W = \bar{U}_\perp^\top U_\perp U_\perp^\top \bar{U}_\perp = \bar{U}_\perp^\top \bar{U}_\perp = I_{n-r}$, so $W \in \text{St}_{\mathbb{R}}(n-r, n-r)$. The second identity in (3.15) becomes $W^\top (U_\perp^\top Z U_\perp) = 0$, i.e., $\text{im}(U_\perp^\top Z U_\perp) \subseteq (\text{im } W)^\perp$ in $\mathbb{R}^{n-r}$. Since W has $n - r$ orthonormal columns, $\text{rank}(W) = n - r$, hence

$$\text{rank}(U_\perp^\top Z U_\perp) \leq \dim(\text{im } W)^\perp = (n - \bar{r}) - (n - r) = r - \bar{r}.$$

Step (II). This argument is the symmetric counterpart of [21, Section 5.2]. Let $Z = Z_1 + Z_2$ with $Z_1 \in T_{S_{\bar{r}}(n)}(X)$ and $Z_2 \in S_{\leq r-\bar{r}}(n)$. Since $S_{\bar{r}}(n)$ is a smooth submanifold of $S(n)$ near X , its tangent cone at X coincides with its tangent space [11, Example 6.8], so there exist $(Y_i)_{i \in \mathbb{N}} \subset S_{\bar{r}}(n)$ with $Y_i \rightarrow X$ and $(t_i)_{i \in \mathbb{N}} \subset (0, \infty)$ with $t_i \rightarrow 0$ such that $(Y_i - X)/t_i \rightarrow Z_1$. Set $X_i := Y_i + t_i Z_2$. Then:

- $X_i \in S(n)$ as a sum of symmetric matrices;
- $\text{rank}(X_i) \leq \text{rank}(Y_i) + \text{rank}(Z_2) \leq \bar{r} + (r - \bar{r}) = r$, so $X_i \in S_{\leq r}(n)$;
- $X_i \rightarrow X$ and $(X_i - X)/t_i = (Y_i - X)/t_i + Z_2 \rightarrow Z_1 + Z_2 = Z$.

Hence $Z \in T_{S_{\leq r}(n)}(X)$. □

Corollary 3.3 (Block characterization).

$$T_{S_{\leq r}(n)}(X) = \left\{ [U \ U_\perp] \begin{bmatrix} A & B \\ B^\top & C \end{bmatrix} [U \ U_\perp]^\top \mid \begin{array}{l} A \in S(\bar{r}), \\ B \in \mathbb{R}^{\bar{r} \times (n-\bar{r})}, \\ C \in S_{\leq r-\bar{r}}(n-\bar{r}) \end{array} \right\}.$$

Proof. Substituting (3.5) and (3.6) in (3.11), an element of $T_{S_{\leq r}(n)}(X)$ has the form

$$[U \ U_\perp] \begin{bmatrix} A & B \\ B^\top & 0 \end{bmatrix} [U \ U_\perp]^\top + [U \ U_\perp] \begin{bmatrix} 0 & 0 \\ 0 & C \end{bmatrix} [U \ U_\perp]^\top = [U \ U_\perp] \begin{bmatrix} A & B \\ B^\top & C \end{bmatrix} [U \ U_\perp]^\top,$$

where the constraint $C \in S_{\leq r-\bar{r}}(n-\bar{r})$ comes from the intersection in (3.11). □

3.2.4 Projection onto the tangent cone

Proposition 3.4. *For $Z \in \mathbb{R}^{n \times n}$, the orthogonal projection of Z onto $T_{S_{\leq r}(n)}(X)$ in $\mathbb{R}^{n \times n}$ with respect to $\|\cdot\|_F$ is*

$$P_{T_{S_{\leq r}(n)}(X)}(Z) = [U \ U_{\perp}] \begin{bmatrix} \hat{A} & \hat{B} \\ \hat{B}^{\top} & \hat{C} \end{bmatrix} [U \ U_{\perp}]^{\top}, \quad (3.16)$$

where, using the block notation (3.3),

$$\hat{A} = \frac{1}{2}(Z_{11} + Z_{11}^{\top}), \quad (3.17)$$

$$\hat{B} = \frac{1}{2}(Z_{12} + Z_{21}^{\top}), \quad (3.18)$$

$$\hat{C} = P_{S_{\leq r-\bar{r}}(n-\bar{r})}\left(\frac{1}{2}(Z_{22} + Z_{22}^{\top})\right). \quad (3.19)$$

The last projection is computed by truncated eigendecomposition: if $M := \frac{1}{2}(Z_{22} + Z_{22}^{\top}) = \sum_{i=1}^{n-\bar{r}} \lambda_i q_i q_i^{\top}$ with $|\lambda_1| \geq |\lambda_2| \geq \dots$, then $\hat{C} = \sum_{i=1}^{r-\bar{r}} \lambda_i q_i q_i^{\top}$.

Proof. By orthogonal invariance of $\|\cdot\|_F$ and Corollary 3.3,

$$\min_{V \in T_{S_{\leq r}(n)}(X)} \|Z - V\|_F^2 = \min_{\substack{A \in S(\bar{r}), B \in \mathbb{R}^{\bar{r} \times (n-\bar{r})} \\ C \in S_{\leq r-\bar{r}}(n-\bar{r})}} \left\| \begin{bmatrix} Z_{11} - A & Z_{12} - B \\ Z_{21} - B^{\top} & Z_{22} - C \end{bmatrix} \right\|_F^2.$$

The squared Frobenius norm splits blockwise, so the problem decouples into three independent subproblems.

Block A. Projection of Z_{11} onto the linear subspace $S(\bar{r})$ yields (3.17).

Block B. The cost over B is $\|Z_{12} - B\|_F^2 + \|Z_{21} - B^{\top}\|_F^2 = \|Z_{12} - B\|_F^2 + \|Z_{21}^{\top} - B\|_F^2$; the first-order condition gives (3.18).

Block C. For symmetric S and skew-symmetric K of the same size, $\langle S, K \rangle_F = \text{tr}(SK) = \text{tr}(K^{\top} S^{\top}) = -\text{tr}(SK) = 0$, so the symmetric and skew-symmetric parts of Z_{22} are $\|\cdot\|_F$ -orthogonal:

$$\|Z_{22} - C\|_F^2 = \left\| \frac{1}{2}(Z_{22} + Z_{22}^{\top}) - C \right\|_F^2 + \left\| \frac{1}{2}(Z_{22} - Z_{22}^{\top}) \right\|_F^2.$$

The optimal C is the projection of $M = \frac{1}{2}(Z_{22} + Z_{22}^{\top})$ onto $S_{\leq r-\bar{r}}(n-\bar{r})$. For symmetric M , the singular value decomposition reduces to the spectral decomposition with singular values $|\lambda_i|$, so the Eckart–Young theorem [25, Theorem 5.5, 5.7] gives (3.19). $\square$

3.3 Order-three real symmetric matrices with bounded rank

Geometry of 3×3 Symmetric Matrices of Rank at Most 1

Following the insights gained from the 2×2 case, the numerical investigation is extended to the set of 3×3 real symmetric matrices of rank at most one, denoted by $S_{\leq 1}(3)$. This part aims to verify if the topological and dynamical patterns identified on the double cone (specifically the impact of the step size α and the presence of singular traps) persist in higher-dimensional settings.

3.3.1 Geometry and Parameterization of $S_{\leq 1}(3)$

As established by [23] and discussed in the previous section, the variety $S_{\leq 1}(n)$ can be partitioned into a smooth manifold $S_1(n)$ and a singular point at the origin $S_0(n) = \{\mathbf{0}\}$. For $n = 3$, the smooth part $S_1(3)$ remains composed of $r + 1 = 2$ connected components, corresponding to the positive-semidefinite (PSD) and negative-semidefinite (NSD) folds.

While the 2×2 case allowed for a linear isomorphism to a three-dimensional double cone, $S(3)$ is a six-dimensional object. To maintain a visual intuition, the following rank-one parameterization is used:

- **PSD component:** $A = uu^\top$ for $u \in \mathbb{R}^3 \setminus \{0\}$.
- **NSD component:** $A = -uu^\top$ for $u \in \mathbb{R}^3 \setminus \{0\}$.

This representation is not an isometry. Specifically, the mapping is 2-to-1 (both u and $-u$ map to the same matrix A) and the Frobenius norm $\|A\|_F$ corresponds to the square of the Euclidean norm $\|u\|^2$. Nevertheless, this representation preserves the critical boundary between the parts: the singularity at $u = 0$.

3.3.2 Numerical Setup and Convergence Targets

To ensure consistency with the 2×2 experiments, the objective function remains the quadratic cost function where the global minimum is a rank-one NSD matrix:

$$A^* = \begin{bmatrix} 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{bmatrix}. \quad (3.20)$$

In our NSD u -space, this target corresponds to the points $u^* = (0, \pm 1, 0)$. As previously defined, an apocalypse is declared if the iterates of Algorithm 2 converge toward the origin, where the projected gradient vanishes regardless of the proximity to A^* .

3.3.3 Results: Persistence of the Step-Size phenomenon

A grid search over the cube $u \in [-5, 5]^3$ is performed with 27000 points to characterize the basin of attraction of the apocalypse. The results for various step sizes are illustrated in Figure 3.4.




Figure 3.4: three-dimensional visualization of the apocalypse basin in u -space. All points indicate initial conditions leading to the singularity (origin). As α decreases, the singularity's basin expands substantially.

The experiments confirm that the step-size phenomenon observed on the double cone is a characteristic of P^2GD on rank-deficient varieties. For large step sizes ($\alpha = 0.7$), the algorithm often possesses a sufficiently large displacement to bypass the singularity's zone of influence. Conversely, as α decreases, the basin of attraction of the apocalypse seems to expand, eventually absorbing nearly the entire feasible domain as soon as $\alpha \leq 0.4$.

3.3.4 Structural Observations

Two structural features emerge:

1. **The stability corridor:** A hole remains visible around the u_2 -axis (the y -axis). This region contains initial points that are naturally aligned with the direction of $A^* = \text{diag}(0, -1, 0)$ (i.e. all matrices of the form $\text{diag}(0, t, 0)$ for all $t \in \mathbb{R}$). While this corridor narrows as $\alpha \rightarrow 0$, it represents the only reliable path to convergence.
2. **The apocalypse plane:** Regardless of the step size, a plane exists at $u_2 = 0$ where all points other than the origin lead to the apocalypse. For any starting point (except the origin) in this plane, the component of the gradient toward the minimum is zero ($\nabla f(A) = A - A^*$) is not visualizable easily in this representation). Consequently, these points necessarily converge toward the origin.

Remark: If the minimum is shifted to $A^* = \text{diag}(-1, 0, 0)$, the stability corridor rotates toward the u_1 -axis and the apocalypse plane shifts to $u_1 = 0$ (and the same goes for u_3).

Diagonal Invariance and Axial Convergence. A property observed in both $S_{\leq 1}(2)$ and $S_{\leq 1}(3)$ concerns the behavior of diagonal initial matrices. Let the global minimum be a diagonal matrix $A^* = \text{diag}(0, \dots, \lambda_i, \dots, 0)$ where only the i -th diagonal element is non-zero and $\lambda_i \in \mathbb{R}$. Our experiments show that:

- Any diagonal initial matrix $A_0 = \text{diag}(0, \dots, a_i, \dots, 0)$ sharing the same non-zero index as the minimum A^* will converge to the global minimum, regardless of the initial value $a_i \in \mathbb{R}$.
- Conversely, any diagonal matrix A_0 whose non-zero element is at an index $j \neq i$ will converge to the apocalypse.

Linking to the analysis in 3.1, it seems that the gradient flow for diagonal matrices remains restricted to the axes. If the initial matrix has no component along the i -th axis, the projected gradient cannot draw the iterates in that direction. This structural constraint would explain why the stability corridor is strictly aligned with the coordinate axes of the parameter space u .

3.4 Higher-dimensional Generalizations

To assess the robustness of the patterns identified in previous sections, this study extends the numerical analysis to higher-dimensional symmetric matrices and

higher-rank varieties. Specifically, the varieties $S_{\leq r}(n)$ for $(n, r) = (5, 2)$ and $(n, r) = (10, 2)$ are investigated. This transition introduces additional complexity, as the singular set $S_{< r}(n)$ is no longer reduced to the origin but contains sub-varieties of intermediate ranks.

3.4.1 Experimental Protocol

The experimental setup is designed to evaluate the probability of convergence toward different stationary points as a function of the step size α . For each dimension $n \in \{5, 10\}$, a target matrix $A^* \in S_r(n)$ is defined as a diagonal matrix of rank r . Initial points X_0 are generated by sampling matrices from the positive-semidefinite (PSD) fold of the variety. Specifically, $X_0 = UU^\top$ is set where $U \in \mathbb{R}^{n \times r}$ is a matrix whose entries are sampled from a standard normal distribution. For each value of $\alpha \in [0.04, 0.7]$, a batch of $N = 200$ simulations is conducted using Algorithm 2 with the following parameters:

- **Step size α :** Fixed value per batch, ranging from 0.1 to 0.7.
- **Backtracking:** Armijo parameters $c = 10^{-4}$ and $\beta = 0.5$.
- **Stopping criteria:** Maximum iterations set to 5000 or a stationarity tolerance $\epsilon = 10^{-4}$.

3.4.2 Classification of Convergence Outcomes

In the context of $S_{\leq r}(n)$ where $r > 1$, an apocalypse can occur at various levels of rank deficiency. The outcome of each trial is classified according to the final iterate X_{final} and its proximity to the global minimum or the singular strata:

$$\text{Outcome} = \begin{cases} \text{Success} & \text{if } \|X_{\text{final}} - A^*\|_F < 10^{-4}, \\ \text{Apocalypse (Rank 0)} & \text{if } \|X_{\text{final}}\|_F < 10^{-4}, \\ \text{Apocalypse (Rank } k) & \text{if } \text{rank}_\tau(X_{\text{final}}) = k \text{ with } 0 < k < r. \end{cases} \quad (3.21)$$

The rank of a matrix A is computed as the number of singular values that exceed a prescribed threshold $\tau = 10^{-4}$ in our case. More precisely, if $\sigma_1(A) \geq \dots \geq \sigma_{\min(m,n)}(A)$ denote the singular values of $A \in \mathbb{R}^{m \times n}$, we set

$$\text{rank}_\tau(A) := \#\{i : \sigma_i(A) > \tau\}.$$

This classification is set to reflect the hierarchical structure of the singular set: the algorithm may be trapped at the origin (rank 0) or on a singular stratum of intermediate rank (rank 1 in this case), where the norm of the projected gradient (i.e. the stationarity measure) goes to zero despite the objective function not being minimized.

3.4.3 Numerical Results and Analysis

The results of the simulations for $n = 5$ and $n = 10$ (both for $r = 2$) are presented in Figure 3.5.




Figure 3.5: Distribution of convergence outcomes for P²GD on $S_{\leq 2}(n)$ as a function of the step size α .

Persistence of the Step-Size phenomenon. The data demonstrates that the correlation between step size and convergence stability remains consistent with observations in lower dimensions. For both $n = 5$ and $n = 10$, small step sizes ($\alpha \leq 0.2$) lead to a failure of the algorithm in most cases, with a 100% probability of convergence toward the rank-0 apocalypse as α decreases.

Emergence of Intermediate Rank Apocalypses. Higher-rank experiments reveal rank-one apocalypses. Unlike the $r = 1$ case where only the origin was a singular trap, the variety $S_{\leq 2}(n)$ contains a singular stratum $S_1(n)$.

- For $\alpha \in [0.04, 0.2]$, a non-negligible proportion of trials (approximately 10-15%) converging to a rank-1 matrix is observed.
- This suggests that the algorithm can be trapped not only at the origin but also on any lower-rank manifold where the tangential geometry prevents further descent.

Convergence Regime for Large Step Sizes. For $\alpha \geq 0.2$, the success rate continues to increase with the step-size and reaches around 96% for $n = 5$ and around 99% for $n = 10$. The remaining 4% (or 1%) of trials predominantly converge

to rank-1 singular points. The proportion of rank-0 apocalypses almost disappear as α increases, suggesting that larger steps enables the iterates to overshoot the most severe singularities, although intermediate traps can still be an obstacle. The similarity between the results for $n = 5$ and $n = 10$ suggests that the dimension of the ambient space has a secondary effect compared to the rank r and the step size α .

Chapter 4

Apocalypse Analysis: a Continuous-Time Perspective

Building on the work of [15, Section 1], this chapter develops a continuous-time surrogate of the projected projected-gradient descent algorithm (P²GD) on the variety $S_{\leq 1}(2)$. Extending the research of [19, Section 4.2], the goal is to explain analytically the apocalyptic behavior observed numerically in Chapter 1. The surrogate is the differential inclusion obtained as the $\alpha \rightarrow 0$ limit of P²GD. On the upper nappe of the double cone, it reduces to a planar autonomous system in cylindrical coordinates (θ, z) .

Main results. We obtain two main contributions. First, in Section 4.3, we construct an iterative scheme of integral bounds on z and $\cos \theta$, yielding a sufficient condition (Corollary 4.5) on the initial angle θ_0 for the trajectory to be trapped near the apex. Second, in Section 4.5, the result is strengthened into a necessary and sufficient statement: for every initial condition $(\theta_0, z_0) \in (0, 2\pi) \times (0, \infty)$, the maximal solution is global in time and satisfies $z(t) > 0$ for all $t \geq 0$ (Theorem 4.7). The trajectories of the continuous model therefore never reach the apex in finite time, which explains why the apocalyptic basin of P²GD covers nearly the entire domain as $\alpha \rightarrow 0$.

4.1 The Projected Projected-Gradient Descent (P²GD)

4.1.1 A Continuous-Time Model for P²GD

Let E be a Euclidean space, $C \subseteq E$ a nonempty closed set, and $f : E \rightarrow \mathbb{R}$ a differentiable function with locally Lipschitz continuous gradient. The projected

projected-gradient descent (P²GD) algorithm 2 selects, at each iterate $x \in C$, a descent direction

$$g(x) \in P_{T_C(x)}(-\nabla f(x)), \quad (4.1)$$

and then updates the iterate by a step followed by a projection onto C :

$$x^+ \in P_C(x + \alpha g(x)), \quad (4.2)$$

where the step size $\alpha > 0$ is chosen by a backtracking line-search satisfying an Armijo-type decrease condition.

To gain analytical insight into the behavior of P²GD, we consider its small-step-size regime $\alpha \rightarrow 0$. Interpreting the update (4.2) as an explicit Euler step for the vector field $x \mapsto g(x)$, one obtains the continuous-time surrogate presented in [19]

$$\dot{x}(t) = P_{T_C(x(t))}(-\nabla f(x(t))). \quad (4.3)$$

On a smooth stratum, for instance, the fixed-rank manifold $S_r(n)$, the set C locally coincides with a $\mathcal{C}^2$ submanifold $M \subseteq E$, in order that $T_C(x) = T_M(x)$ is a linear subspace. In this case, the metric projection P_M is locally single-valued and differentiable around any $\bar{x} \in M$, with derivative $DP_M(\bar{x}) = P_{T_M(\bar{x})}$ ([26, Lemma 4]). Therefore, for any $u \in T_M(x)$ and for $\alpha > 0$ sufficiently small, a first-order Taylor expansion of P_M at x yields

$$P_M(x + \alpha u) = P_M(x) + \alpha DP_M(x)[u] + o(\alpha) = x + \alpha P_{T_M(x)}(u) + o(\alpha) = x + \alpha u + o(\alpha).$$

Applied to the P²GD iteration

$$x_{k+1} = P_M(x_k + \alpha g(x_k)), \quad g(x) = P_{T_C(x)}(-\nabla f(x)),$$

and noting that $g(x_k) \in T_M(x_k)$ on the smooth stratum, this shows that

$$x_{k+1} = x_k + \alpha g(x_k) + o(\alpha),$$

in order that the discrete-time update is a first-order (explicit Euler) discretization of the continuous-time dynamics (4.3).

At a singular point of C , such as the origin of the double cone, the tangent cone $T_C(x)$ is in general nonconvex and the projection $P_{T_C(x)}(\cdot)$ may be set-valued. In this case, (4.3) is more appropriately interpreted as the differential inclusion

$$\dot{x}(t) \in P_{T_C(x(t))}(-\nabla f(x(t))).$$

4.1.2 The Projected Gradient Flow on the Double Cone

This section specializes the continuous-time dynamics (4.3) to the setting analyzed numerically in the previous chapter. Specifically, the variety $S_{\leq 1}(2)$ is locally identified with the double cone C , defined as:

$$C := \{x \in \mathbb{R}^3 : x_3^2 = x_1^2 + x_2^2\}. \quad (4.4)$$

The quadratic objective function $f(x) = \frac{1}{2}\|x - x^*\|^2$ with the specific target point $x^* = (-1, 0, -1)^\top$ as presented in Chapter 3 is considered. The negative gradient of f is given by:

$$-\nabla f(x) = x^* - x = \begin{bmatrix} -1 - x_1 \\ -x_2 \\ -1 - x_3 \end{bmatrix}. \quad (4.5)$$

The continuous-time P²GD dynamics is defined as the projection of this negative gradient onto the tangent space $T_C(x)$, expressed as $\dot{x} = P_{T_C(x)}(-\nabla f(x))$. Using the relation $P_{T_C(x)}(g) = g - P_{N_C(x)}(g)$, the projection onto the normal space

$$P_{N_C(x)}(-\nabla f(x)) = \frac{\langle n(x), x^* - x \rangle}{\|n(x)\|^2} n(x). \quad (4.6)$$

Since $x \in C$, it follows that $n(x) = (x_1, x_2, -x_3)^\top$ and $\|n(x)\|^2 = x_1^2 + x_2^2 + x_3^2 = 2x_3^2$. The inner product yields:

$$\begin{aligned} \langle n(x), x^* - x \rangle &= x_1(-1 - x_1) + x_2(-x_2) + (-x_3)(-1 - x_3) \\ &= -x_1 - (x_1^2 + x_2^2) + x_3 + x_3^2 \\ &= x_3 - x_1, \end{aligned}$$

where the terms x_3^2 and $-(x_1^2 + x_2^2)$ cancel due to the cone constraint. Thus, the dynamics in Cartesian coordinates is:

$$\dot{x} = \begin{bmatrix} -1 - x_1 \\ -x_2 \\ -1 - x_3 \end{bmatrix} - \frac{x_3 - x_1}{2x_3^2} \begin{bmatrix} x_1 \\ x_2 \\ -x_3 \end{bmatrix}. \quad (4.7)$$

On the upper nappe $C_+ = \{x \in C : x_3 > 0\}$, the cone constraint $x_3^2 = x_1^2 + x_2^2$ implies $x_3 = \sqrt{x_1^2 + x_2^2}$, so a single non-negative coordinate $z := x_3$ doubles as the cylindrical radius. We therefore parameterize C_+ by

$$x_1 = z \cos \theta, \quad x_2 = z \sin \theta, \quad x_3 = z, \quad z \in (0, +\infty), \quad \theta \in [0, 2\pi). \quad (4.8)$$

Note that the target $x^* = (-1, 0, -1)^\top$ lies on the *lower* nappe, so every trajectory considered here is, by construction, separated from the global minimum by the

apex—the question is whether it reaches the apex. First, the vertical coordinate z satisfies $\dot{z} = \dot{x}_3$:

$$\begin{aligned}\dot{z} &= (-1 - z) - \frac{z - z \cos \theta}{2z^2}(-z) \\ &= -1 - z + \frac{1 - \cos \theta}{2} \\ &= -z - \cos^2 \frac{\theta}{2},\end{aligned}$$

using the trigonometric identity $\frac{1 - \cos \theta}{2} = 1 - \cos^2 \frac{\theta}{2}$. Next, for the angular coordinate θ , using the relation $z\dot{\theta} = \dot{x}_2 \cos \theta - \dot{x}_1 \sin \theta$

$$\begin{aligned}z\dot{\theta} &= \left(-z \sin \theta - \frac{z - z \cos \theta}{2z^2}z \sin \theta\right) \cos \theta - \left(-1 - z \cos \theta - \frac{z - z \cos \theta}{2z^2}z \cos \theta\right) \sin \theta \\ &= -z \sin \theta \cos \theta + \sin \theta + z \cos \theta \sin \theta \\ &= \sin \theta,\end{aligned}$$

where the projection terms $\frac{1 - \cos \theta}{2z}(x_2 \cos \theta - x_1 \sin \theta)$ vanish because $x_2 \cos \theta - x_1 \sin \theta = z \sin \theta \cos \theta - z \cos \theta \sin \theta = 0$. Consequently, the dynamical system reduces to the following planar ODE found in [19]:

$$\begin{cases} \dot{\theta} = \frac{\sin \theta}{z} \\ \dot{z} = -z - \cos^2 \frac{\theta}{2} \end{cases}, \quad z \in (0, +\infty) \text{ and } \theta \in (0, 2\pi) \quad (4.9)$$

Existence and uniqueness of a maximal solution. Define $U := (0, 2\pi) \times (0, +\infty)$ and

$$F(\theta, z) = \left(\frac{\sin \theta}{z}, -z - \cos^2 \left(\frac{\theta}{2} \right) \right).$$

Since $z > 0$ on U , the map $(\theta, z) \mapsto 1/z$ is smooth on U , and therefore $F \in C^\infty(U, \mathbb{R}^2)$. In particular, F is locally Lipschitz on U . By the Cauchy–Lipschitz theorem ([27, Chapter 6]), for every $(\theta_0, z_0) \in U$ there exists a unique maximal solution

$$(\theta, z) : I_{\max} \rightarrow U, \quad I_{\max} = (t_-, t_+), \quad (4.10)$$

such that $(\theta(0), z(0)) = (\theta_0, z_0)$.

Two qualitative properties of (4.9) are immediate. First, since $\frac{d}{dt} \cos \theta(t) = -\sin^2 \theta(t)/z(t) \leq 0$, the quantity $\cos \theta(t)$ is monotonically nonincreasing along any trajectory. Second, the inequality $\dot{z} = -z - \cos^2(\theta/2) \leq -z$ implies $z(t) \leq z_0 e^{-t}$, so $z(t)$ decays at least exponentially. The central question, and the subject of the remainder of this chapter, is whether $z(t)$ can reach zero in finite time (i.e. escape the singular point).

4.2 Exact Integral Equations

The analysis of (4.9) rests on two exact integral representations, which transform the differential system into a form amenable to comparison arguments.

Proposition 4.1 (Exact integral equations). *Let (z, θ) be the maximal solution of (4.9) with initial data $z_0 > 0$ and $\theta_0 \in (0, 2\pi) \setminus \{\pi\}$. On its interval of existence, the following identities hold:*

$$z(t) = z_0 e^{-t} - e^{-t} \int_0^t e^s \cos^2 \frac{\theta(s)}{2} ds, \quad (4.11)$$

$$\cos \theta(t) = \tanh \left(\operatorname{artanh}(\cos \theta(0)) - \int_0^t \frac{ds}{z(s)} \right). \quad (4.12)$$

Proof. For (4.11): the ODE $\dot{z} = -z - \cos^2(\theta/2)$ is linear in z . Multiplying both sides by the integrating factor e^t gives

$$\dot{z}e^t + ze^t = \frac{d}{dt}(ze^t) = -e^t \cos^2(\theta/2);$$

integrating from 0 to t and rearranging yields

$$z(t)e^t - z_0 = - \int_0^t e^s \cos^2 \frac{\theta(s)}{2} ds$$

dividing the expression by e^t results in (4.11).

For (4.12): compute

$$\frac{d}{dt} \cos \theta = -\frac{\sin^2 \theta}{z} = -\frac{1 - \cos^2 \theta}{z}.$$

Separating variables and integrating from 0 to t

$$\int_{\cos \theta(0)}^{\cos \theta(t)} \frac{dx}{1 - x^2} = - \int_0^t \frac{ds}{z(s)}$$

using $\int \frac{dc}{1-c^2} = \operatorname{artanh}(c)$ yields

$$\operatorname{artanh}(\cos \theta(t)) = \operatorname{artanh}(\cos \theta(0)) - \int_0^t \frac{ds}{z(s)}$$

Applying $\tanh$ on both sides results in (4.12). □

4.3 Bounds on the Trajectory

Proposition 4.1 cannot be used directly to bound $z(t)$ and $\cos \theta(t)$, since each integral depends on the other unknown. To evaluate whether the system reaches the origin in finite time (which would signify the capacity to bypass the apocalypse), we must rigorously quantify the trajectory's behavior near the singularity. We therefore construct monotone sequences of upper and lower bounds by alternately refining estimates on z and $\cos \theta$, using (4.11)–(4.12) at each step.

4.3.1 Iterative Refinement of the Bounds

Starting from obvious initial bounds, an improvement can be systematically obtained by iterating the coupling between z and $\cos \theta$. Let $c_L^{(n)}, c_U^{(n)}$ denote lower and upper bounds on $\cos \theta(t)$ at step n , and let $z_L^{(n)}, z_U^{(n)}$ denote the corresponding bounds on $z(t)$. Initialize

$$z_L^{(0)}(t) := -\infty, \quad z_U^{(0)}(t) := +\infty, \quad c_L^{(0)}(t) := -1, \quad c_U^{(0)}(t) := 1.$$

For $n \geq 0$, define the refined bounds by

$$z_L^{(n+1)}(t) := z_0 e^{-t} - e^{-t} \int_0^t e^s \frac{1 + c_U^{(n)}(s)}{2} ds, \quad (\text{B}_1)$$

$$z_U^{(n+1)}(t) := z_0 e^{-t} - e^{-t} \int_0^t e^s \frac{1 + c_L^{(n)}(s)}{2} ds, \quad (\text{B}_2)$$

$$c_L^{(n+1)}(t) := \tanh \left(\operatorname{artanh}(\cos \theta_0) - \int_0^t \frac{ds}{z_L^{(n+1)}(s)} \right), \quad (\text{B}_3)$$

$$c_U^{(n+1)}(t) := \tanh \left(\operatorname{artanh}(\cos \theta_0) - \int_0^t \frac{ds}{z_U^{(n+1)}(s)} \right). \quad (\text{B}_4)$$

The rationale is as follows. A tighter upper bound $c_U^{(n)}(t)$ on $\cos \theta$ yields a tighter (larger) lower bound $z_L^{(n+1)}(t)$ on z via (B₁), since a smaller $\cos^2(\theta/2)$ reduces the integral subtracted from $z_0 e^{-t}$. Conversely, a larger $z_L^{(n+1)}(t)$ makes $1/z_L^{(n+1)}(t)$ smaller, reducing the integral in (B₃) and yielding a tighter lower bound on $\cos \theta$.

Proposition 4.2 (Monotone enclosure). *For all $n \geq 0$ and $t \geq 0$,*

$$z_L^{(n)}(t) \leq z(t) \leq z_U^{(n)}(t), \quad c_L^{(n)}(t) \leq \cos \theta(t) \leq c_U^{(n)}(t).$$

Proof. Proceeding by induction on n . The base case is immediate for $z_U^{(0)}(t) \geq z(t) \geq z_L^{(0)}(t)$. The bounds on $\cos \theta(t)$, $c_L^{(0)}(t) = -1 \leq \cos \theta(t) \leq 1 = c_U^{(0)}(t)$ holds trivially.

Assume the inequalities hold at step n . We prove them at step $n + 1$.

Lower bound on z ((B₁)). Using the identity $\cos^2(\theta/2) = (1 + \cos \theta)/2$ and the induction hypothesis $\cos \theta(s) \leq c_U^{(n)}(s)$, one has

$$\cos^2 \frac{\theta(s)}{2} = \frac{1 + \cos \theta(s)}{2} \leq \frac{1 + c_U^{(n)}(s)}{2}.$$

Substituting into (4.11) yields $z(t) \geq z_L^{(n+1)}(t)$.

Upper bound on z ((B₂)). By the induction hypothesis $\cos \theta(s) \geq c_L^{(n)}(s)$, the integrand $(1 + \cos \theta(s))/2 \geq (1 + c_L^{(n)}(s))/2$, so the subtracted integral in (4.11) is larger, giving $z(t) \leq z_U^{(n+1)}(t)$.

Lower bound on $\cos \theta$ (B3). Suppose that $z_L^{(n+1)}(s) > 0$ for all $s \in [0, t]$ (this holds, for instance, under the sufficient condition of Corollary 4.5). From the lower bound on z just established, $z(s) \geq z_L^{(n+1)}(s) > 0$, hence $1/z(s) \leq 1/z_L^{(n+1)}(s)$. Substituting into (4.12):

$$\cos \theta(t) \geq c_L^{(n+1)}(t) = \tanh \left(\operatorname{artanh}(\cos \theta_0) - \int_0^t \frac{ds}{z_L^{(n+1)}(s)} \right).$$

Upper bound on $\cos \theta$ (B4). Since $z_U^{(n+1)}(s) \geq z(s) > 0$ by the induction hypothesis and the positivity of z , we have $1/z(s) \geq 1/z_U^{(n+1)}(s) > 0$. Substituting into (4.12):

$$\cos \theta(t) \leq c_U^{(n+1)}(t) = \tanh \left(\operatorname{artanh}(\cos \theta_0) - \int_0^t \frac{ds}{z_U^{(n+1)}(s)} \right).$$

□

4.3.2 First Level Bounds

The level $n = 1$ gives the first nontrivial bounds.

Lower bound on z : Substituting $c_U^{(0)}(t) = 1$ in (B₁) yields

$$\begin{aligned} z_L^{(1)}(t) &= z_0 e^{-t} - e^{-t} \int_0^t e^s ds \\ &= z_0 e^{-t} - e^{-t}(e^t - 1) \\ &= e^{-t}(z_0 + 1) - 1. \end{aligned}$$

Upper bound on z : Substituting $c_L^{(0)}(t) = -1$ in (B₂) yields

$$\begin{aligned} z_U^{(1)}(t) &= z_0 e^{-t} - e^{-t} \int_0^t 0 \, ds \\ &= z_0 e^{-t}. \end{aligned}$$

Lower bound on $\cos \theta$: Substituting $z_L^{(0)}(t) = e^{-t}(z_0 + 1) - 1$ in (B₃) yields

$$c_L^{(1)}(t) = \tanh \left(\operatorname{artanh}(\cos \theta_0) - \int_0^t \frac{ds}{e^{-t}(z_0 + 1) - 1} \right).$$

A change of variable is performed using $u = e^{-s}(z_0 + 1)$ and $ds = -\frac{du}{u}$ the integral becomes

$$\int_0^t \frac{ds}{e^{-t}(z_0 + 1) - 1} = \int_{e^{-t}(z_0+1)}^{z_0+1} \frac{du}{u(u-1)} = \int_{e^{-t}(z_0+1)}^{z_0+1} \left(\frac{1}{u-1} - \frac{1}{u} \right) du.$$

Solving the integral gives

$$\left[\ln \left| \frac{u-1}{u} \right| \right]_{e^{-t}(z_0+1)}^{z_0+1} = \ln \left| \frac{z_0}{z_0 + 1 - e^t} \right|.$$

Substituting in the initial equation yields

$$c_L^{(1)}(t) = \tanh \left(\operatorname{artanh}(\cos \theta_0) - \ln \left| \frac{z_0}{z_0 + 1 - e^t} \right| \right).$$

Upper bound on $\cos \theta$: Substituting $z_U^{(0)}(t) = z_0 e^{-t}$ in (B₄) yields

$$c_U^{(1)}(t) = \tanh \left(\operatorname{artanh}(\cos \theta_0) - \int_0^t \frac{ds}{z_0 e^{-t}} \right).$$

Solving the integral yields

$$c_U^{(1)}(t) = \tanh \left(\operatorname{artanh}(\cos \theta_0) - \frac{e^t - 1}{z_0} \right).$$

This expression is equivalent to the expression found in [19, Proposition 4]:

$$c_U^{(1)}(t) = \frac{AE(t) - 1}{AE(t) + 1}, \quad (4.13)$$

where

$$A := \frac{1 + \cos \theta_0}{1 - \cos \theta_0}, \quad E(t) := \exp\left(2 \frac{1 - e^t}{z_0}\right). \quad (4.14)$$

Note that $A > 0$ for $\theta_0 \in (0, \pi)$, and $E(t) \rightarrow 0$ as $t \rightarrow +\infty$, so the right-hand side of (4.13) converges to -1 from above, consistently with $\theta(t) \rightarrow \pi$.

Proposition 4.3. *The two following expressions for the upper bound of the cosine of the angle, $c_U^{(1)}(t)$, are equivalent:*

$$c_U^{(1)}(t) = \tanh\left(\operatorname{artanh}(\cos \theta_0) - \frac{e^t - 1}{z_0}\right) = \frac{AE(t) - 1}{AE(t) + 1} \quad (4.15)$$

where $A := \frac{1 + \cos \theta_0}{1 - \cos \theta_0}$ and $E(t) := \exp\left(2 \frac{1 - e^t}{z_0}\right)$.

Proof. Using the subtraction formula for the hyperbolic tangent, $\tanh(a - b) = \frac{\tanh(a) - \tanh(b)}{1 - \tanh(a)\tanh(b)}$, let $a = \operatorname{artanh}(\cos \theta_0)$ and $b = \frac{e^t - 1}{z_0}$. Noting that $\tanh(a) = \cos \theta_0$:

$$c_U^{(1)}(t) = \frac{\cos \theta_0 - \tanh(b)}{1 - \cos \theta_0 \tanh(b)}$$

Substituting $\tanh(b) = \frac{1 - E(t)}{1 + E(t)}$ into the equation using the exponential definition of $\tanh$:

$$c_U^{(1)}(t) = \frac{\cos \theta_0 - \frac{1 - E(t)}{1 + E(t)}}{1 - \cos \theta_0 \left(\frac{1 - E(t)}{1 + E(t)}\right)} = \frac{\cos \theta_0(1 + E(t)) - (1 - E(t))}{(1 + E(t)) - \cos \theta_0(1 - E(t))}$$

Rearranging terms with respect to $E(t)$ yields:

$$c_U^{(1)}(t) = \frac{E(t)(\cos \theta_0 + 1) - (1 - \cos \theta_0)}{E(t)(\cos \theta_0 + 1) + (1 - \cos \theta_0)}$$

Dividing the numerator and denominator by $(1 - \cos \theta_0)$ leads to the final form:

$$c_U^{(1)}(t) = \frac{E(t) \left(\frac{1 + \cos \theta_0}{1 - \cos \theta_0}\right) - 1}{E(t) \left(\frac{1 + \cos \theta_0}{1 - \cos \theta_0}\right) + 1} = \frac{AE(t) - 1}{AE(t) + 1}$$

□

4.3.3 Refined Lower Bound on $z(t)$

Using $c_U^{(1)}(t)$ the upper bound on $\cos \theta(s)$ (4.13) yields

$$\frac{1 + c_U^{(1)}(t)}{2} = \frac{AE(s)}{AE(s) + 1} =: H(s). \quad (4.16)$$

Evaluating the integral via the substitution $u = E(s)$, which gives $du = -\frac{2e^s}{z_0} E(s) ds$, one obtains

$$\int_0^t \frac{AE(s)}{AE(s) + 1} e^s ds = -\frac{z_0}{2} \int_{E(0)}^{E(t)} \frac{A}{Au + 1} du = \frac{z_0}{2} \ln \left(\frac{A + 1}{AE(t) + 1} \right).$$

This yields the following explicit lower bound.

Proposition 4.4 (Lower bound on $z(t)$). *Under the notation of (4.14), for all $t \geq 0$,*

$$z(t) \geq z_L^{(2)}(t) = z_0 e^{-t} \left[1 - \frac{1}{2} \ln \left(\frac{A + 1}{AE(t) + 1} \right) \right]. \quad (4.17)$$

Proof. Substituting the upper bound (4.16) into (B₁):

$$z_L^{(2)}(t) = z_0 e^{-t} - e^{-t} \int_0^t e^s \frac{1 + c_U^{(1)}}{2} ds = z_0 e^{-t} - e^{-t} \int_0^t H(s) e^s ds.$$

The explicit evaluation of the integral above then yields (4.17). $\square$

4.3.4 Refined Lower Bound on $\cos \theta$

Substituting the lower bound (4.17) for $z(s)$ in (4.12) (which increases the integral since the denominator is now smaller) gives a lower bound on $\cos \theta(t)$ involving $z_{\min}(s)$:

$$\cos \theta(t) \geq \tanh \left(\operatorname{artanh}(\cos \theta_0) - \int_0^t \frac{ds}{z_{\min}(s)} \right), \quad (4.18)$$

where $z_{\min}(s)$ denotes the right-hand side of (4.17). The integral in (4.18) can be expressed in closed form via the substitution $u = E(s)$:

$$\int_0^t \frac{ds}{z_{\min}(s)} = \int_0^t \frac{e^s}{z_0 \left[1 - \frac{1}{2} \ln \left(\frac{A+1}{AE(s)+1} \right) \right]} ds = \frac{1}{2} \int_{E(t)}^1 \frac{1}{u \left[1 - \frac{1}{2} \ln \left(\frac{A+1}{Au+1} \right) \right]} du. \quad (4.19)$$

4.3.5 Initial-Angle Criterion

The first-level lower bound $z_L^{(1)}$ provides the first rigorous analytical certificate of global existence for the P²GD continuous-time model. This bound allows to identify a specific set of initial conditions (z_0, θ_0) for which $z(t)$ is guaranteed to remain strictly positive for all $t \geq 0$, (i.e. remaining trapped by the apocalypse). Since $E(t) \rightarrow 0$ as $t \rightarrow +\infty$, the logarithmic term in (4.17) converges to $\ln(A + 1)$. The lower bound on $z(t)$ therefore remains positive for all $t \geq 0$ whenever

$$1 - \frac{1}{2} \ln(A + 1) \geq 0 \iff A + 1 \leq e^2. \quad (4.20)$$

To express this in terms of the initial angle θ_0 , the half-angle identity $\cot(\theta_0/2) = \sqrt{(1 + \cos \theta_0)/(1 - \cos \theta_0)}$ is used, which gives $A = \cot^2(\theta_0/2)$. Substituting into (4.20):

$$\cot^2 \frac{\theta_0}{2} + 1 < e^2 \iff \frac{1}{\sin^2(\theta_0/2)} < e^2 \iff \sin \frac{\theta_0}{2} > \frac{1}{e}.$$

Corollary 4.5 (Sufficient condition for apocalyptic trapping). *Let $(\theta_0, z_0) \in ((0, 2\pi) \setminus \{\pi\}) \times (0, +\infty)$. If*

$$2 \arcsin\left(\frac{1}{e}\right) < \theta_0 < 2\pi - 2 \arcsin\left(\frac{1}{e}\right), \quad (4.21)$$

then $z_L^{(2)}(t) > 0$ for all $t \geq 0$, and consequently $z(t) \geq z_L^{(2)}(t) > 0$. The trajectory therefore does not reach the apex $z = 0$ in finite time, and the maximal solution is global.

The threshold $2 \arcsin(1/e) \approx 0.753 \text{ rad} \approx 43.17^\circ$ delimits two regimes: initial angles sufficiently far from 0 and 2π are guaranteed to remain trapped by (2.17), while angles close to these values may, a priori, escape. The numerical experiments of Chapter 1 are consistent with this picture: the basin of attraction of the apocalypse was found to concentrate around $\theta = 0$ as $\alpha \rightarrow 0$. Theorem 4.7 below will show that, in fact, *every* non-degenerate initial condition is trapped, so (4.21) is sufficient but far from necessary.

4.4 Numerical Illustration of the Iterative Scheme

Analytical refinement of the bounds beyond level $n = 1$ becomes intractable, owing to the nesting of the integrals and the transcendental form of the dynamics. We therefore evaluate the higher-order bounds numerically and assess the convergence of the iterative scheme (B₁)–(B₄).

Validity interval of the bounds. As established in Section 4.3, the lower bound $c_L^{(n+1)}$ on $\cos \theta$ derived from (B₃) is valid only on intervals where $z_L^{(n+1)}(t) > 0$. Outside this regime $z_L^{(n+1)}$ remains a valid lower bound on z , while $c_L^{(n+1)}$ ceases to be guaranteed. To ensure that the two-sided enclosure is well defined across all iterations, we define the validity horizon

$$T := \sup \left\{ t \geq 0 : z_L^{(n)}(s) > 0 \text{ for all } s \in [0, t] \text{ and all } n \in \{0, \dots, 4\} \right\},$$

and restrict the nested enclosure to $[0, T]$, on which the bounds at every iteration $n \leq 4$ are simultaneously valid.

Numerical setup. The exact solution $(z(t), \theta(t))$ of (4.9) is computed by adaptive Runge–Kutta integration (RK45) with absolute tolerance 10^{-10} . The bounds $z_L^{(n)}, z_U^{(n)}, c_L^{(n)}, c_U^{(n)}$ are evaluated by computing the integrals (B₁)–(B₄) with the trapezoidal rule on a uniform grid of 10^4 points. The scheme is iterated up to $n = 4$. We fix $z_0 = 1$ and consider three initial angles,

$$\theta_0 \in \{\pi/10, \pi/5, \pi/4\} \approx \{0.31, 0.63, 0.79\},$$

two of which lie below and one above the threshold $\theta_\star := 2 \arcsin(1/e) \approx 0.75$ of Corollary 4.5.

Convergence of the scheme. Figure 4.1 displays the successive enclosures of $z(t)$ and $\cos \theta(t)$ over $[0, T]$ for $\theta_0 = \pi/10$ and $\theta_0 = \pi/5$. The bands $[z_L^{(n)}, z_U^{(n)}]$ and $[c_L^{(n)}, c_U^{(n)}]$ contract around the exact solution at each iteration, the tightest enclosure ($n = 4$) coinciding visually with the RK45 reference curve. Figure 4.2 quantifies this contraction: the enclosure width $\max_{t \in [0, T]} (z_U^{(n)}(t) - z_L^{(n)}(t))$ (and its analogue for $\cos \theta$) decreases geometrically with n . After four iterations the maximum uncertainty on z falls by roughly four orders of magnitude (from $\sim 10^1$ to $\sim 10^{-3}$), and that on $\cos \theta$ by three orders of magnitude.

Pointwise-best bounds on a long horizon. The enclosure of Figure 4.1 is confined to the short common horizon $[0, T]$, on which every order $n \leq 4$ is simultaneously valid. This horizon is too short to expose the asymptotic behaviour of z , and in particular whether z stays positive as t grows. To reach the long-time regime we drop the requirement that a single order be valid throughout and retain, at each time t , the tightest bound among the orders still valid at t . With $N = 4$ fixed, define the envelope bounds

$$\widehat{z}_L(t) := \max \left\{ z_L^{(n)}(t) : 0 \leq n \leq N, z_L^{(n)} \text{ valid on } [0, t] \right\}, \quad \widehat{z}_U(t) := \min_{0 \leq n \leq N} z_U^{(n)}(t),$$




Figure 4.1: Successive enclosures of $z(t)$ (left) and $\cos \theta(t)$ (right) produced by the iterative scheme (B_1) – (B_4) , restricted to the validity horizon $[0, T]$. Top: $\theta_0 = \pi/5 \approx 0.63$. Bottom: $\theta_0 = \pi/10 \approx 0.31$. Solid bands denote the enclosure at iteration n ; the dashed black curve is the RK45 reference solution.




Figure 4.2: Maximum width of the enclosure (log scale) for z and $c = \cos \theta$, measured on $[0, T]$, as a function of the iteration index n , for $(z_0, \theta_0) = (1, \pi/10)$.

where “ $z_L^{(n)}$ valid on $[0, t]$ ” means that the positivity condition underlying (B_3) has not failed on $[0, t]$ for any order feeding the recursion. By Proposition 4.2, each retained bound satisfies $z_L^{(n)}(t) \leq z(t) \leq z_U^{(n)}(t)$, hence $\hat{z}_L(t) \leq z(t) \leq \hat{z}_U(t)$ for all $t \in \mathbb{R}_+$.

Two elementary members of the family govern the envelope at large t . The order-one upper bound, obtained from $c_L^{(0)} = -1$, is

$$z_U^{(1)}(t) = z_0 e^{-t};$$


it carries no positivity condition, is valid for all $t \geq 0$, and forces $\hat{z}_U(t) \leq z_0 e^{-t}$, so the upper envelope stays positive and decays to zero on the whole horizon. The order-one lower bound, obtained from $c_U^{(0)} = 1$, is

$$z_L^{(1)}(t) = (z_0 + 1) e^{-t} - 1,$$


also valid for all t , but positive only for $t < \ln(z_0 + 1)$. Past this time the lower envelope can be held above zero only by higher-order iterates, which remain valid exactly when the sufficient condition of Corollary 4.5 is met.

Behaviour across the critical angle. Figure 4.3 reads directly through the threshold θ_* . For $\theta_0 = \pi/4 > \theta_*$, the sufficient condition of Corollary 4.5 holds, the higher-order lower bounds stay valid for all t , and the two envelopes remain glued together and strictly positive over the entire horizon $[0, 5]$; the enclosure is consistent with $z(t) > 0$ at every computed time. For $\theta_0 = \pi/5$ and $\theta_0 = \pi/10$, both below θ_* , the envelopes coincide until the order-one lower bound loses positivity near $t = \ln(z_0 + 1) = \ln 2 \approx 0.69$; past this time the lower envelope detaches, dips below zero, and then recovers towards zero as t grows, the forcing $\cos^2(\theta/2)$ vanishing as $\theta \rightarrow \pi$, while the upper envelope keeps tracking the decay $z_0 e^{-t}$. The dip is deeper for the smaller angle ($\hat{z}_L \approx -0.2$ at $\theta_0 = \pi/10$ against ≈ -0.04 at $\theta_0 = \pi/5$), which quantifies the degradation of the lower bound as θ_0 moves further below the threshold.


Interpretation. The envelopes separate what the elementary bound machinery can and cannot certify. Above the critical angle the lower envelope stays positive, so the order-four enclosure certifies $z(t) > 0$ on the full horizon, independently of Theorem 4.7. Below it the certificate is only partial: the region of uncertainty contains zero past $t = \ln(z_0 + 1)$, so the lower envelope alone does not exclude $z = 0$ even though it looks like z is going to stay strictly positive. That exclusion is supplied by Theorem 4.7, which guarantees $z(t) > 0$ for every initial condition; the upper envelope, positive and decaying as $z_0 e^{-t}$, is consistent with this and confines z to a shrinking neighbourhood of the apex it never reaches.



(a) $\theta_0 = \pi/10 \approx 0.31$ (below θ_*).



(b) $\theta_0 = \pi/5 \approx 0.63$ (below θ_*).



(c) $\theta_0 = \pi/4 \approx 0.79$ (above θ_*).

Figure 4.3: Envelope bounds $\hat{z}_L$ and $\hat{z}_U$ on $z(t)$ over $[0, 5]$, obtained at budget $N = 4$ by retaining the tightest valid bound at each time, with $z_0 = 1$ and three initial angles. The threshold is $\theta_* = 2 \arcsin(1/e) \approx 0.75$. The shaded band is the region of uncertainty $[\hat{z}_L, \hat{z}_U]$.

4.5 Global existence

The analytical bounds derived in Section 2.3 certify global existence only for a subset of initial conditions. Numerical experiments, however, consistently suggest that no trajectory of the continuous-time system reaches the singular apex $z = 0$ in finite time. We prove this fact here: for every initial condition $(\theta_0, z_0) \in U$, the maximal solution is defined for all $t \geq 0$ and satisfies $z(t) > 0$.

4.5.1 Continuous-time system and basic properties

We recall the autonomous system on $U = (0, 2\pi) \times (0, +\infty)$:

$$\begin{cases} \dot{\theta} = \frac{\sin \theta}{z}, \\ \dot{z} = -z - \cos^2\left(\frac{\theta}{2}\right), \end{cases} \quad (\theta(0), z(0)) = (\theta_0, z_0) \in U. \quad (4.22)$$

By (4.10), for every $(\theta_0, z_0) \in U$ there exists a unique maximal solution $(\theta, z) : I_{\max} \rightarrow U$, with $I_{\max} = (t_-, t_+)$ and $(\theta(0), z(0)) = (\theta_0, z_0)$.

The line $\mathcal{M} := \{(\theta, z) \in U : \theta = \pi\}$ is forward invariant. Indeed, if $\theta_0 = \pi$, the system reduces to $\dot{\theta} = 0$ and $\dot{z} = -z$, yielding the explicit global solution $\theta(t) \equiv \pi$ and $z(t) = z_0 e^{-t} > 0$ for $t \geq 0$. By uniqueness, no trajectory with $\theta_0 \neq \pi$ can meet $\mathcal{M}$. The involution $\theta \mapsto 2\pi - \theta$ further preserves (4.22). We may therefore restrict, without loss of generality, to $\theta_0 \in (0, \pi)$. The invariance of $\mathcal{M}$ then forces $\theta(t) \in (0, \pi)$ for all $t \in [0, t_+)$, and consequently

$$\dot{\theta}(t) = \frac{\sin \theta(t)}{z(t)} > 0, \quad \dot{z}(t) = -z(t) - \cos^2(\theta(t)/2) < 0,$$

so θ is strictly increasing and z is strictly decreasing on $[0, t_+)$.

4.5.2 Global existence ($t_+ = \infty$)

We first record a trigonometric inequality used in the proof of the main theorem.

Lemma 4.6. *For every $\epsilon \in (0, \pi]$ and every $x \in [0, \epsilon]$,*

$$\sin x \geq \left(1 - \frac{\epsilon^2}{6}\right) x.$$

Proof. The inequality holds at $x = 0$. For $x > 0$, it is equivalent to $\sin(x)/x \geq 1 - \epsilon^2/6$. The map $f(x) := \sin(x)/x$ is strictly decreasing on $(0, \pi]$: indeed $f'(x) = (x \cos x - \sin x)/x^2$, and $g(x) := x \cos x - \sin x$ satisfies $g(0) = 0$ and

$g'(x) = -x \sin x \leq 0$ on $[0, \pi]$, so $g \leq 0$ and $f' \leq 0$. For $x \in (0, \epsilon]$ we therefore have $\sin(x)/x \geq \sin(\epsilon)/\epsilon$, and it suffices to prove

$$\sin \epsilon \geq \epsilon - \frac{\epsilon^3}{6}.$$

By Taylor's theorem applied to $\sin$ at 0 with Lagrange remainder of order 4, there exists $\xi \in (0, \epsilon)$ such that

$$\sin \epsilon = \epsilon - \frac{\epsilon^3}{6} + \frac{\sin \xi}{24} \epsilon^4.$$

Since $\epsilon \leq \pi$, $\xi \in (0, \pi)$, hence $\sin \xi > 0$ and the inequality follows. $\square$

Theorem 4.7. *For every initial condition $(\theta_0, z_0) \in U$, the maximal time of existence satisfies $t_+ = +\infty$ and $z(t) > 0$ for all $t \geq 0$.*

Proof. By the monotonicity established above, the limits

$$\theta_* := \lim_{t \rightarrow t_+^-} \theta(t) \in (\theta_0, \pi], \quad \ell := \lim_{t \rightarrow t_+^-} z(t) \in [0, z_0)$$

both exist. By the escape-from-compact property of maximal solutions [28, Cor.2.16], it suffices to show that the trajectory is contained in a compact subset of U . We distinguish two cases according to the value of θ_* .

Case 1: $\theta_* \in (\theta_0, \pi)$. We show that z is bounded below by a positive constant on $[0, t_+)$; the trajectory then remains in a compact subset of U , and the escape-from-compact property yields $t_+ = +\infty$.

Since θ is monotonically nondecreasing, $\theta(t) \in [\theta_0, \theta_*]$ for every $t \in [0, t_+)$, and

$$m := \min_{\theta \in [\theta_0, \theta_*]} \sin \theta > 0$$

because $[\theta_0, \theta_*] \subset (0, \pi)$.

Since $z \in C^1([0, t_+))$ and $\dot{z}(t) < 0$ for all $t \in [0, t_+)$, the map z is strictly decreasing, hence bijective onto $z([0, t_+))$. By the C^1 inverse function theorem in one dimension ([29, Theorem 9.24]), z^{-1} is C^1 on $z([0, t_+))$ with $(z^{-1})'(s) = 1/z'(z^{-1}(s))$. Denote $z^{-1} : z([0, t_+)) \rightarrow [0, t_+)$ and set $\tilde{\theta} := \theta \circ z^{-1}$. The chain rule, together with $z(z^{-1}(s)) = s$, gives for every $s \in z([0, t_+))$

$$-\tilde{\theta}'(s) = -\frac{\theta'(z^{-1}(s))}{z'(z^{-1}(s))} = \frac{\sin(\theta(z^{-1}(s)))}{s(s + \cos^2(\theta(z^{-1}(s))/2))}.$$

Using $\sin(\theta(z^{-1}(s))) \geq m$, $\cos^2(\cdot) \leq 1$ and $s \leq z_0$,

$$-\tilde{\theta}'(s) \geq \frac{m}{s(z_0 + 1)}.$$

Integrating from $z(t)$ to z_0 yields, for every $t \in [0, t_+)$,

$$\theta(t) - \theta_0 = - \int_{z(t)}^{z_0} \tilde{\theta}'(s) ds \geq \frac{m}{z_0 + 1} \ln \left(\frac{z_0}{z(t)} \right). \quad (4.23)$$

Inequality (4.23) embodies the following alternative: along the trajectory, $z(t)$ cannot become arbitrarily small without $\theta(t)$ growing arbitrarily large. Since $\theta(t) \leq \theta_* < \pi$ on $[0, t_+)$, the left-hand side of (4.23) is bounded by $\theta_* - \theta_0$, which forces

$$\ln \left(\frac{z_0}{z(t)} \right) \leq \frac{(\theta_* - \theta_0)(z_0 + 1)}{m} =: C, \quad \text{i.e.} \quad z(t) \geq z_0 e^{-C} =: \delta_1 > 0.$$

In particular $\ell \geq \delta_1 > 0$. The trajectory is therefore contained in the compact set

$$K_1 := [\theta_0, \theta_*] \times [\delta_1, z_0] \subset U.$$

It yields that if $\theta_* \in (\theta_0, \pi)$ then $t_+ = +\infty$.

Case 2: $\theta_* = \pi$. Assume, for contradiction, that $t_+ < +\infty$. By the escape-from-compact property and $\theta_* = \pi \in (0, 2\pi)$, this forces $\ell = 0$.

Set $x(t) := \pi - \theta(t)$, so $x(t) > 0$ on $[0, t_+)$ and $x(t) \rightarrow 0$ as $t \rightarrow t_+^-$. System (4.22) becomes

$$\dot{x} = -\frac{\sin x}{z}, \quad \dot{z} = -z - \sin^2 \left(\frac{x}{2} \right).$$

Fix $\epsilon \in (0, \sqrt{3})$ and set $\alpha := 1 - \epsilon^2/6 \in (1/2, 1)$. Choose $t_1 \in [0, t_+)$ such that $x(t) \leq \epsilon$ for all $t \in [t_1, t_+)$.

As in Case 1, z is a C^1 -diffeomorphism, and we set $\tilde{x} := x \circ z^{-1}$ on $z([t_1, t_+))$. The chain rule and $z(z^{-1}(s)) = s$ yield

$$\tilde{x}'(s) = \frac{x'(z^{-1}(s))}{z'(z^{-1}(s))} = \frac{\sin(\tilde{x}(s))}{s(s + \sin^2(\tilde{x}(s)/2))}.$$

By Lemma 4.6, $\sin(\tilde{x}(s)) \geq \alpha \tilde{x}(s)$, and since $\sin^2(\tilde{x}/2) \leq 1$,

$$\frac{\tilde{x}'(s)}{\tilde{x}(s)} \geq \frac{\alpha}{s(s + 1)} = \alpha \left(\frac{1}{s} - \frac{1}{s + 1} \right).$$

Integrating from $z(t)$ to $z(t_1)$ gives

$$\ln \left(\frac{x(t_1)}{x(t)} \right) \geq \alpha \ln \left(\frac{z(t_1)}{z(t_1) + 1} \cdot \frac{z(t) + 1}{z(t)} \right),$$

and exponentiating, together with $z/(z+1) \leq z$ for $z \geq 0$,

$$x(t) \leq K z(t)^\alpha, \quad K := x(t_1) \left(\frac{z(t_1) + 1}{z(t_1)} \right)^\alpha.$$

Since $\sin u \leq u$ for $u \geq 0$,

$$\sin^2\left(\frac{x(t)}{2}\right) \leq \frac{x(t)^2}{4} \leq \frac{K^2}{4} z(t)^{2\alpha}.$$

Because $2\alpha > 1$, there exists $\delta > 0$ such that $\frac{K^2}{4} z^{2\alpha} \leq z$ for every $z \in (0, \delta]$. As $z(t) \rightarrow \ell = 0$, we may choose $t_2 \in [t_1, t_+)$ with $z(t) \leq \delta$ for all $t \in [t_2, t_+)$. For such t ,

$$\dot{z}(t) = -z(t) - \sin^2(x(t)/2) \geq -2z(t),$$

Thus, multiplying by e^{2t} gives

$$\frac{d}{dt}(e^{2t} z(t)) = e^{2t}(\dot{z}(t) + 2z(t)) \geq 0,$$

so $t \mapsto e^{2t} z(t)$ is monotonically nondecreasing on $[t_2, t_+)$ and

$$z(t) \geq z(t_2) e^{-2(t-t_2)} \geq z(t_2) e^{-2(t_+-t_2)} > 0 \quad \text{for all } t \in [t_2, t_+).$$

The trajectory on $[t_2, t_+)$ is therefore contained in the compact set

$$K_2 := [\theta(t_2), \pi] \times [z(t_2) e^{-2(t_+-t_2)}, z(t_2)] \subset U,$$

contradicting the escape-from-compact property under the assumption $t_+ < +\infty$. Hence $t_+ = +\infty$ in this case as well.

In both cases $t_+ = +\infty$, and consequently $z(t) > 0$ for all $t \geq 0$. $\square$

4.5.3 Implication for the discrete algorithm

Theorem 4.7 concerns the continuous-time surrogate (4.9), which was obtained in Section 4.1 as the small-step limit of P²GD: for $\alpha \rightarrow 0$, the iteration $x^+ \in P_C(x + \alpha g(x))$ is a first-order (explicit Euler) discretization of (4.9). To the extent that the discrete iterates track this flow at small step size, the continuous picture transfers to the algorithm. Since the continuous trajectory never reaches the apex (Theorem 4.7), it approaches the apex ever more closely—over an horizon that grows as $\alpha \rightarrow 0$ —without ever escaping the upper nappe. This is consistent with the central numerical observation of Chapter 3: the apocalyptic basin under P²GD fills nearly the entire upper nappe as $\alpha \rightarrow 0$, the sole exception being trajectories initialised on the descent line $x_2 = 0$, $x_1 > 0$, which lead to the rank-one minimizer.

Conclusion

This thesis asked whether the apocalypse afflicting the projected projected-gradient descent method (P^2GD) is a marginal pathology, tolerable in practice, or a structural feature of low-rank optimization that genuinely warrants the corrective schemes proposed in the literature. We addressed the question on the variety of symmetric matrices of bounded rank, whose geometry is explicit enough to support both quantitative experiments and a closed-form continuous-time analysis. The evidence gathered along the way points consistently to the second answer: the apocalypse is structural.

On the geometric side, we adapted the variational analysis of the rectangular case to the symmetric variety, obtaining closed-form expressions for its tangent cone and for the orthogonal projection onto it (Theorem 3.2, Proposition 3.4). Beyond their role in making P^2GD implementable in arbitrary dimension, these formulas isolate the mechanism behind the apocalypse: as smooth iterates approach a low-rank point, their tangent spaces omit directions of decrease that the tangent cone at the singular point itself retains. Along the sequence, the projected anti-gradient becomes increasingly orthogonal to these narrower feasible directions, so the stationarity measure collapses toward zero; yet the singular point is not stationary—a descent direction persists there—and the algorithm wrongly reports convergence.

On the experimental side, we mapped the set of initial conditions driven to an apocalypse, from the double cone up to the higher-dimensional varieties of rank at most two in dimensions five and ten. Two regularities stand out. First, the step size influences the phenomenon: as it decreases, the apocalyptic set grows until it dominates almost the entire feasible domain, while the basin of the global minimizer shrinks to a single descent line; at the smallest steps tested, the great majority of trajectories are trapped by a singular point. Second, and more subtly, increasing the step size does not cure the phenomenon uniformly across the singular strata. Larger steps allow the iterates to bypass the origin, but once the singular set is stratified a residual fraction of trajectories settles instead on an intermediate-rank stratum, where the tangential geometry again prevents further descent. Nevertheless, the proportion of intermediate-rank apocalypses decreases substantially as the step

size grows past a threshold.

On the theoretical side, we analyzed the continuous-time limit of P^2GD , its projected gradient flow, which on the double cone decouples into scalar equations for the height and the angle. We constructed a sequence of iteratively refined upper and lower bounds on the trajectory (Proposition 4.2) and extracted from them an explicit sufficient condition on the initial angle for the iterates to be driven toward the apex (Corollary 4.5). Numerically, these enclosures contract at a geometric rate: after a handful of iterations the certified band is visually indistinguishable from a high-accuracy reference solution, and it does so even for initial angles that lie outside the regime where the sufficient condition is guaranteed, confirming that the practical reach of the scheme exceeds what the analysis alone certifies. We then proved that the flow is globally defined for every initial condition and that its height remains strictly positive for all finite time while tending to zero asymptotically (Theorem 4.7), settling a question left open in [19]. The continuous trajectory is thus trapped: it approaches the singular point without ever reaching it, and never turns back toward the minimizer.

These three elements together give a coherent picture. The trapping established in the continuous-time system explains the central empirical observation, that the apocalyptic basin fills the domain as the step size tends to zero. Although a larger step size appears to reduce the probability of an apocalypse, it is not sufficient on its own to guarantee that the phenomenon is eliminated. The apocalypse is, in this sense, intrinsic to non-smooth varieties such as the bounded-rank varieties, rather than an accident of a particular implementation, and the corrective schemes of the literature address a real problem.

These conclusions are nonetheless limited by the scope of this study. The continuous-time analysis is carried out only on the double cone; the higher-dimensional varieties are explored numerically, without an analogous closed-form flow. The experiments use a single quadratic objective with a prescribed minimizer, which isolates the role of the geometry but leaves the dependence on the cost function open. Finally, the analysis lives in the idealized small-step limit, whereas practical P^2GD uses an Armijo line search, and the two are linked here only to first order.

These limitations chart the work that remains. The main direction of interest is to carry the analysis back from the continuous flow to the discrete algorithm itself: to describe its convergence directly at the level of the iterates, rather than through the continuous surrogate: delineating the apocalypse basin, characterizing precisely which initial conditions and step sizes lead to failure, and turning that characterization into a practical criterion. In particular, it would be valuable to determine theoretically whether a sufficiently large step size can preclude apocalyptic behaviour. Another natural question is the role of the objective: replacing the

quadratic cost by a broader family of functions would test whether the step-size dependence and the intermediate-rank traps are features of the geometry or artifacts of this particular case. Lifting the closed-form analysis to higher-dimensions, above the double cone, would, in turn, test whether the intermediate-rank apocalypses observed numerically admit the same exact description as the rank-zero trap. Beyond P^2GD , the apocalypse is shared by other feasible methods on bounded-rank varieties, such as the retraction-free descent (RFD) of [15, Algorithm 4 and Sec 3.4]; determining whether the trapping mechanism identified here transfers to these algorithms, and whether a single geometric remedy addresses them all, would clarify how general the phenomenon really is.

References

- [1] E. Candès and B. Recht, “Exact matrix completion via convex optimization,” *Commun. ACM*, vol. 55, no. 6, pp. 111–119, Jun. 2012, ISSN: 0001-0782. DOI: 10.1145/2184319.2184343. [Online]. Available: <https://doi.org/10.1145/2184319.2184343>.
- [2] J. Tanner and K. Wei, “Normalized iterative hard thresholding for matrix completion,” *SIAM Journal on Scientific Computing*, vol. 35, no. 5, S104–S125, 2013. DOI: 10.1137/120876459. eprint: <https://doi.org/10.1137/120876459>. [Online]. Available: <https://doi.org/10.1137/120876459>.
- [3] B. Recht, M. Fazel, and P. A. Parrilo, “Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization,” *SIAM Review*, vol. 52, no. 3, pp. 471–501, 2010. DOI: 10.1137/070697835. eprint: <https://doi.org/10.1137/070697835>. [Online]. Available: <https://doi.org/10.1137/070697835>.
- [4] Y. Chi, Y. M. Lu, and Y. Chen, “Nonconvex optimization meets low-rank matrix factorization: An overview,” *IEEE Transactions on Signal Processing*, vol. 67, no. 20, pp. 5239–5269, 2019. DOI: 10.1109/TSP.2019.2937282.
- [5] E. J. Candès, X. Li, Y. Ma, and J. Wright, “Robust principal component analysis?” *J. ACM*, vol. 58, no. 3, Jun. 2011, ISSN: 0004-5411. DOI: 10.1145/1970392.1970395. [Online]. Available: <https://doi.org/10.1145/1970392.1970395>.
- [6] J. Abernethy, F. Bach, T. Evgeniou, and J.-P. Vert, “A new approach to collaborative filtering: Operator estimation with spectral regularization,” *Journal of Machine Learning Research*, vol. 10, no. 29, pp. 803–826, 2009. [Online]. Available: <http://jmlr.org/papers/v10/abernethy09a.html>.
- [7] N. Gillis and F. Glineur, “Low-rank matrix approximation with weights or missing data is NP-hard,” *SIAM J. Matrix Anal. Appl.*, vol. 32, no. 4, pp. 1149–1165, 2011. DOI: <https://doi.org/10.1137/110820361>.

- [8] N. Srebro and T. Jaakkola, “Weighted low-rank approximations,” in *Proceedings of the 20th international conference on machine learning (ICML-03)*, 2003, pp. 720–727.
- [9] X. Zhou, C. Yang, H. Zhao, and W. Yu, “Low-rank modeling and its applications in image analysis,” *ACM Comput. Surv.*, vol. 47, no. 2, Dec. 2014, ISSN: 0360-0300. DOI: 10.1145/2674559. [Online]. Available: <https://doi.org/10.1145/2674559>.
- [10] S. Schotthöfer, E. Zangrando, J. Kusch, G. Ceruti, and F. Tudisco, “Low-rank lottery tickets: Finding efficient low-rank neural networks via matrix differential equations,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 20 051–20 063, 2022. DOI: <https://doi.org/10.48550/arXiv.2205.13571>.
- [11] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis* (Grundlehren der mathematischen Wissenschaften). Berlin: Springer, 1998, vol. 317, ISBN: 978-3-540-62772-2. DOI: 10.1007/978-3-642-02431-3.
- [12] G. Olikier and I. Waldspurger, “Projected gradient descent accumulates at bouligand stationary points,” *SIAM Journal on Optimization*, vol. 35, no. 2, pp. 1004–1029, 2025. DOI: 10.1137/24M1692782.
- [13] A. Uschmajew and B. Vandereycken, “Geometric methods on low-rank matrix and tensor manifolds,” in *Handbook of Variational Methods for Nonlinear Geometric Data*, P. Grohs, M. Holler, and A. Weinmann, Eds., Cham: Springer International Publishing, 2020, pp. 261–313, ISBN: 978-3-030-31351-7. DOI: 10.1007/978-3-030-31351-7_9. [Online]. Available: https://doi.org/10.1007/978-3-030-31351-7_9.
- [14] K. Wei, J.-F. Cai, T. F. Chan, and S. Leung, “Guarantees of riemannian optimization for low rank matrix recovery,” *SIAM Journal on Matrix Analysis and Applications*, vol. 37, no. 3, pp. 1198–1222, 2016. DOI: 10.1137/15M1050525. eprint: <https://doi.org/10.1137/15M1050525>. [Online]. Available: <https://doi.org/10.1137/15M1050525>.
- [15] R. Schneider and A. Uschmajew, “Convergence results for projected line-search methods on varieties of low-rank matrices via Łojasiewicz inequality,” *SIAM J. Optim.*, vol. 25, no. 1, pp. 622–646, 2015. DOI: 10.1137/140957822.
- [16] G. Olikier, K. A. Gallivan, and P.-A. Absil, “First-order optimization on stratified sets,” Preprint, Mar. 2023. DOI: 10.48550/arXiv.2303.16040.

- [17] E. Levin, J. Kileel, and N. Boumal, “Finding stationary points on bounded-rank matrices: A geometric hurdle and a smooth remedy,” *Mathematical Programming*, vol. 199, no. 1, pp. 831–864, May 2023, ISSN: 1436-4646. DOI: 10.1007/s10107-022-01851-2. [Online]. Available: <https://doi.org/10.1007/s10107-022-01851-2>.
- [18] G. Olikier, K. A. Gallivan, and P.-A. Absil, “Low-rank optimization methods based on projected projected-gradient descent that accumulate at Bouligand stationary points,” *Math. Oper. Res.*, vol. to appear, 2026. DOI: <https://doi.org/10.1287/moor.2024.0582>.
- [19] G. Olikier, “Stability of the apocalypses generated by p2gd on the set of order-two real symmetric singular matrices,” Unpublished manuscript, 2024.
- [20] N. Boumal, *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, 2023, ISBN: 978-1-009-16616-4. DOI: 10.1017/9781009166164.
- [21] G. Olikier, P. Mlinarić, P.-A. Absil, and A. Uschmajew, “The tangent cone to the real determinantal variety: Various expressions and a proof,” *Set-Valued Var. Anal.*, vol. 34, no. 8, 2026. DOI: 10.1007/s11228-026-00796-4.
- [22] E. Arbarello, M. Cornalba, P. A. Griffiths, and J. Harris, *Geometry of Algebraic Curves* (Grundlehren der mathematischen Wissenschaften). New York: Springer, 1985, vol. 267. DOI: 10.1007/978-1-4757-5323-3.
- [23] U. Helmke and J. B. Moore, *Optimization and Dynamical Systems* (Comm. Control Engrg.). London: Springer, 1994, 3rd printing 1996, ISBN: 978-1-4471-3469-5. DOI: 10.1007/978-1-4471-3467-1.
- [24] J. M. Lee, *Introduction to Smooth Manifolds* (Grad. Texts in Math.), 2nd. New York, NY: Springer, 2013, vol. 218, ISBN: 978-1-4419-9981-8. DOI: 10.1007/978-1-4419-9982-5.
- [25] L. N. Trefethen and D. Bau III, *Numerical Linear Algebra*. Philadelphia, PA: SIAM, 1997. DOI: 10.1137/1.9780898719574. [Online]. Available: <https://epubs.siam.org/doi/abs/10.1137/1.9780898719574>.
- [26] P.-A. Absil and J. Malick, “Projection-like retractions on matrix manifolds,” *SIAM J. Optim.*, vol. 22, no. 1, pp. 135–158, 2012. DOI: 10.1137/100802529.
- [27] S. H. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*, 2nd. Boulder, CO: Westview Press, 2015, ISBN: 978-0-8133-4910-7. DOI: 10.1201/9780429492563.
- [28] G. Teschl, *Ordinary differential equations and dynamical systems*. American Mathematical Soc., 2012, vol. 140. DOI: 10.1090/gsm/140.

- [29] W. Rudin, *Principles of Mathematical Analysis*. New York: McGraw-Hill, 1976, vol. 3. DOI: <https://doi.org/10.1017/S0013091500008889>.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain

Rue Archimède, 1 bte L6.11.01, 1348 Louvain-la-Neuve, Belgique | www.uclouvain.be/epl