

École polytechnique de Louvain

Sum-of-Squares programming via optimization on manifolds

Author: **Charles VAN HEES**

Supervisors: **Prof. Pierre-Antoine ABSIL, Prof. Benoît LEGAT**

Reader: **Prof. Nick VANNIEUWENHOVEN**

Academic year 2025–2026

Master [120] in Mathematical Engineering

Abstract

For decades, the nonnegativity of polynomials has been of concern to mathematicians across several fields: global optimization, control theory, probability, game theory, operations research, etc. Nevertheless, in the general case, determining whether a polynomial is nonnegative or not is a NP-hard problem. For univariate polynomials however, nonnegativity is equivalent to being a sum of two squares: there exist two polynomials such that the sum of their squares equals the given polynomial.

A few years ago, Legat et al. derived an algorithm for finding a sum-of-squares decomposition of a univariate polynomial satisfying a set of linear constraints. The novel contribution of this Master's thesis is the introduction of a linear objective we aim to minimize. This problem is traditionally modelled as an SDP and solved using, for instance, an Interior Point Method. However, SDPs quickly become intractable due to the number of variables, which grows quadratically in the size of the problem. Hence, we rely on an idea proposed by Burer and Monteiro at the beginning of the 21st century, by considering an equivalent nonlinear formulation whose variables are the two polynomials of the sum-of-squares decomposition.

Among the various methods that exist to solve such problems (Augmented Lagrangian, Projected Gradient, etc.), we opted for an approach based on manifolds. To this end, we study the feasible set in depth, in order to understand its structure. This allows us to derive a first-order Riemannian optimization algorithm, when the feasible set is a smooth manifold. The algorithm of Legat et al. is then used to obtain a feasible point from which the iterates will start. Although the feasible set is surprisingly smooth in many cases (when the constraint operator is bijective, or near pairs of polynomials sharing at most one common root), it turns out that it is not a smooth manifold for every set of linear constraints defining it. We demonstrate this by showing that the tangent cone is not a vector space at certain points of the set, for a specific set of constraints, around a pair of polynomials sharing two common roots. On such non-smooth sets, further work remains to be done, though we offer some insight by showing that they form stratified sets.

Acknowledgments

I would like to start by thanking my two supervisors, Prof. Pierre-Antoine Absil and Prof. Benoît Legat, for their support throughout the year, and beforehand. Despite the many wrong directions I have taken, they always took the time to advise me and suggesting me new approaches to tackle the problems we encountered. Their expertise in polynomial optimization and manifolds complemented one another, and has proved extremely useful.

Second, I am grateful to Professor Nick Vannieuwenhoven for agreeing to come to UCLouvain to assess my work. Your feedback will certainly be invaluable in advancing the research.

I further thank Katerina Kontogeorgou for taking the time to review my Master's thesis and for providing feedback on the English grammar, spelling and writing style.

I also thank all the people I have encountered in my life who have helped me grow. Whether they guided me in mathematics, Greek, Latin, or Dutch, supported my passion for piano and tennis, or simply inspired me by being them, I will always remember them.

Next, a special thanks to my family for their support, confidence and love over the years. They have believed in me in every project I have undertaken, supporting me whenever I needed help. As the French singer Grand Corps Malade says in his song *Le sens de la famille*, “Ils m’ont transmis tout ce que j’aimerais transmettre à mon tour; c’est grâce à eux si je suis en paix et que je ne pars pas en vrille”¹.

Finally, I thank my friends, with whom I have shared so many lovely moments. Spending five years (or even more for some) together leaves lasting memories I will not forget in a hurry. Grand Corps Malade concludes his song with “La famille, c’est aussi ceux qui sont devenus essentiels; ceux qui te connaissent, te révèlent, te soutiennent et te protègent”².

Use of AI

For my Master's thesis, I made use of artificial intelligence tools (Claude, ChatGPT and DeePL) for three main purposes. The first was to get references on the topics I was studying, mainly for the motivation part. In each case, I analyzed the given references to verify whether they match what I wanted to cover. Secondly, I asked AI to solve nonlinear equations I was facing, to see whether I was going in a right direction. After analyzing the results of the AI (whether they seemed promising or not), I decided to solve the equations myself to get a better understanding (and often in a different way than the AI). Finally, I used AI to check the spelling and grammatical structure of the text, and to improve certain passages at the level of the wording, each time double-checking it myself. Yet, all the work has been done by myself, with the help of my supervisors.

¹Which can be translated as “They have passed on to me everything I would like to pass on in turn; it is thanks to them that I am at peace and haven't lost my way”.

²Which can be translated as “Family also includes those who have become indispensable; those who know you, help you discover who you are, support you and protect you”.

Contents

List of notations	iv
1 Introduction	1
1.1 A few motivating examples	1
1.1.1 Global optimization	1
1.1.2 Existence of Lyapunov functions	2
1.1.3 Equilibria of polynomial games	3
1.2 Sum-of-squares relaxation	3
1.3 Main problem	4
1.4 A Riemannian approach	5
1.5 Structure of the work	7
1.6 Contribution	7
2 Nonnegative and Sum-of-Squares polynomials	8
2.1 Nonnegative and sum-of-squares polynomials	8
2.2 Basis for univariate polynomials	11
2.3 Linear Sum-of-Squares Programming	12
2.4 Nonlinear program for the feasibility problem	12
3 Coprime polynomials as an open submanifold	14
3.1 Topology on pairs of polynomials	14
3.2 Manifold of coprime polynomials	15
4 Manifold structures on the feasible set of sum-of-squares problem	18
4.1 Notations	18
4.2 Embedded submanifold with bijective operator	19
4.3 Embedded submanifolds of coprime polynomials	23
5 Stratification structure on SOS decompositions of polynomials	24
5.1 A failure for our candidate local defining function	24

5.1.1	Communicating vessels	25
5.1.2	A counterexample	26
5.2	Tangent cone	27
5.2.1	Expression of the tangent cone	28
5.2.2	Comments on the tangent cone	29
5.2.3	Tangent cone of smooth embedded submanifolds	30
5.2.4	Tangent cone of our counterexample	31
5.3	A stratified structure	33
5.3.1	Stratification on our example	34
5.3.2	A remedy when $\mathbf{b} = \mathbf{0}$	35
6	First order algorithm for sum-of-squares problem	37
6.1	Riemannian gradient	38
6.1.1	Definition	38
6.1.2	Riemannian gradient for polynomials	39
6.2	Retraction	41
6.2.1	Newton-like retraction	42
	Conclusion	44
	Bibliography	46

List of notations

Mathematical sets

$\mathbb{R}, \mathbb{C}$	Sets of real and complex numbers
$\mathbb{N}, \mathbb{N}_0$	Set of natural numbers, with the zero included and excluded respectively
$\mathbb{R}^m, \mathbb{R}^{m \times n}$	Sets of m -dimensional vectors and $m \times n$ matrices in $\mathbb{R}$
$\mathbb{R}[x]_d$	Set of univariate degree- d polynomials in the variable x , with coefficients in $\mathbb{R}$
$\mathbb{R}[x]_d^r$	Set of r -tuples of univariate polynomials in $\mathbb{R}[x]_d$
$\mathbb{R}[x]_d^{r,c}$	Set of r -tuples of coprime univariate polynomials in $\mathbb{R}[x]_d$
$\mathbb{R}[x_1, \dots, x_n]_d$	Set of degree- d multivariate polynomials with coefficients in $\mathbb{R}$
$\mathbb{S}_+^n$	Set of $n \times n$ symmetric positive semidefinite matrices

Linear Algebra

$\mathcal{E}, \mathcal{F}$	Euclidean spaces
$\ \cdot\ $	Euclidean norm
$\langle \cdot, \cdot \rangle$	Inner product
$\langle \cdot, \cdot \rangle_{\mathcal{E}}$	Inner product on the Euclidean space $\mathcal{E}$ (when precision is needed)
$\mathcal{L} : \mathcal{E} \rightarrow \mathcal{F}$	Linear operator
$\mathcal{L}^* : \mathcal{F} \rightarrow \mathcal{E}$	Adjoint operator of $\mathcal{L} : \mathcal{E} \rightarrow \mathcal{F}$
$\mathcal{L}^{-1}$	Preimage operator of $\mathcal{L} : \mathcal{E} \rightarrow \mathcal{F}$ (returns a subset of $\mathcal{E}$)
$\ker \mathcal{L}, \text{Im } \mathcal{L}$	Kernel and range of the linear operator $\mathcal{L} : \mathcal{E} \rightarrow \mathcal{F}$
$\ker A, \text{Col } A$	Kernel and column space of the matrix $A \in \mathbb{R}^{m \times n}$
$A^\dagger$	Moore-Penrose pseudo-inverse of $A \in \mathbb{R}^{m \times n}$
$\det(A) = A $	Determinant of the square matrix A
$\bar{z}$	Complex conjugate of $z \in \mathbb{C}$
$R(\cdot)$	Rotation matrix
$\mathcal{Z}^\perp$	Orthogonal complement to the linear subspace $\mathcal{Z}$ in $\mathcal{E}$

Polynomials

$p, q \in \mathbb{R}[x]_d$	Univariate polynomials of degree d
$\mathbf{u}, \mathbf{v}, \mathbf{w} \in \mathbb{R}[x]_d^r$	r -tuples of univariate polynomials of degree d
$u_i, v_i, w_i \in \mathbb{R}[x]_d$	Components of $\mathbf{u}, \mathbf{v}, \mathbf{w}$

$K \in \mathbb{N}$	Number of common roots between a set of polynomials
$\sigma : \mathbb{R}[x]_d^r \rightarrow \mathbb{R}[x]_{2d}$	Sum-of-squares operator, defined by $\sigma(\mathbf{u}) = \sum_{i=1}^r u_i^2$
$\mathcal{A}_{\mathbf{u}} : \mathbb{R}[x]_d^r \rightarrow \mathbb{R}[x]_{2d}$	Linear map, defined by $\mathcal{A}_{\mathbf{u}}(\mathbf{v}) = \sum_{i=1}^r u_i v_i$
$\gcd(p, q)$	Greatest common divisor of the polynomials p and q
$\text{Syl}(p, q)$	Sylvester matrix associated to the polynomials p and q
$\text{Res}(p, q)$	Resultant between the polynomials p and q

Riemannian optimization

$\mathcal{M}, \mathcal{N}$	Smooth manifolds
$\mathcal{S}$	Feasible set to the main problem (1.2)
$\mathcal{U}, \mathcal{V}$	Open sets
$\bar{\mathcal{X}}$	Closure of the set $\mathcal{X}$
$h : \mathcal{U} \rightarrow \mathbb{R}^k$	(Candidate) local defining function
$\mathcal{T}_x \mathcal{M}, \mathcal{N}_x \mathcal{M}$	Tangent and normal spaces of $\mathcal{M}$ at $x \in \mathcal{M}$
$0_x \in \mathcal{T}_x \mathcal{M}$	Null tangent vector at $x \in \mathcal{M}$
$\mathcal{T}\mathcal{M}$	Tangent bundle of $\mathcal{M}$
$\langle \cdot, \cdot \rangle_x$	Inner product on the tangent space to a manifold $\mathcal{M}$ at $x \in \mathcal{M}$
$Df(x)[\cdot] : \mathcal{E} \rightarrow \mathcal{F}$	Differential of the function $f : \mathcal{U} \rightarrow \mathcal{V}$ at $x \in \mathcal{U}$
$\nabla f(x)$	Euclidean gradient of the function $f : \mathcal{E} \rightarrow \mathbb{R}$ at $x \in \mathcal{E}$
$Jf(x)$	Jacobian of the function $f : \mathcal{E} \rightarrow \mathbb{R}^k$ at $x \in \mathcal{E}$
$\text{grad } f(x)$	Riemannian gradient of the function $f : \mathcal{M} \rightarrow \mathbb{R}$ at $x \in \mathcal{M}$
$R : \mathcal{T}\mathcal{M} \rightarrow \mathcal{M}$	Retraction on the manifold $\mathcal{M}$
$\text{Proj}_x(\cdot)$	Orthogonal projector from $\mathcal{E}$ to $\mathcal{T}_x \mathcal{M}$, where $\mathcal{M}$ is an embedded submanifold of $\mathcal{E}$
$N_f : \mathcal{E} \rightarrow \mathcal{E}$	Newton operator on the function $f : \mathcal{E} \rightarrow \mathcal{F}$

Abbreviations

SOS	Sum of squares
SDP	Semidefinite programming
gcd	greatest common divisor

Chapter 1

Introduction

Polynomials are deceptively familiar mathematical objects introduced at secondary school. Yet, many of their properties inspire some of the deepest questions in mathematics. In this Master's thesis, we consider polynomials on the field of $\mathbb{R}$, i.e., with coefficients in the reals. A *polynomial* on $\mathbb{R}$ over a set of variables $x_1, \dots, x_n$ is a finite linear combination of monomials with coefficients in $\mathbb{R}$, where a *monomial* over $x_1, \dots, x_n$ is a term of the form $x_1^{\alpha_1} \cdots x_n^{\alpha_n}$, for some n -tuple $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n) \in \mathbb{N}^n$. Bold letters are used here to denote vectors. The *degree* of a monomial corresponds to the sum of the α_i , and the *degree* of a polynomial is the maximal degree of its monomials.

We say that a polynomial is *univariate* if it has only one variable (i.e., $n = 1$). The set of polynomials in $x_1, \dots, x_n$ with coefficients in $\mathbb{R}$ and of degree up to $d \in \mathbb{N}$ is written as $\mathbb{R}[x_1, \dots, x_n]_d$. In particular, for univariate polynomials, we usually omit the subscript 1, and write $\mathbb{R}[x]_d$.

1.1 A few motivating examples

For many applications (operation research, probability theory, etc. [BPT12; AM16]), we are concerned by a specific class of polynomials: *nonnegative polynomials*. A polynomial $p \in \mathbb{R}[x_1, \dots, x_n]_d$ is said to be *nonnegative* if it is nonnegative at every point $(x_1, \dots, x_n) \in \mathbb{R}^n$.

This Master's thesis is concerned with optimization problems on univariate nonnegative polynomials. Before going further, let us start by reviewing three examples (some of them with multivariate polynomials) in global optimization, control theory and game theory.

1.1.1 Global optimization

First, consider the problem of minimizing some polynomial $p(\boldsymbol{x})$ on $\mathbb{R}^n$ [BPT12]:

$$\min_{\boldsymbol{x} \in \mathbb{R}^n} p(\boldsymbol{x}).$$

Many methods have been developed for such problems, like the gradient and Newton methods, and are explained in nonlinear optimization books such as [Nes18]. For the

purpose of the example, consider the equivalent reformulation of the problem as:

$$\begin{aligned} \max_{\gamma \in \mathbb{R}} \quad & \gamma \\ \text{s.t.} \quad & p(\mathbf{x}) - \gamma \geq 0 \quad \forall \mathbf{x} \in \mathbb{R}^n. \end{aligned}$$

This corresponds to finding the best lower bound γ^* on the polynomial $p(\mathbf{x})$. This example may be extended to optimization of rational functions, i.e., functions that can be written as the quotient of two polynomials, in which the denominator is a strictly positive polynomial:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{p(\mathbf{x})}{q(\mathbf{x})} \quad \Longleftrightarrow \quad \begin{aligned} \max_{\gamma \in \mathbb{R}} \quad & \gamma \\ \text{s.t.} \quad & p(\mathbf{x}) - \gamma q(\mathbf{x}) \geq 0 \quad \forall \mathbf{x} \in \mathbb{R}^n. \end{aligned}$$

In these two problems, we seek to maximize γ such that some polynomial is nonnegative. It is interesting to observe that the optimal value γ^* is equal to the global minimum of the polynomial $p(\mathbf{x})$.

1.1.2 Existence of Lyapunov functions

A second example comes from the control community. Indeed, sum-of-squares programming is used in control theory for various applications, including the inverted pendulum, the cart-pole, the quadrotor [Yan+23], robust feedback motion planning [Ani16], etc.

Consider a theoretical system with state $\mathbf{x} \in \mathbb{R}^n$, and described by the equations

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t)).$$

Assuming the system has an equilibrium at the origin, the system is said to be *globally asymptotically stable*¹ [Kha22; BPT12] if there exists a Lyapunov function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ that satisfies for all $\mathbf{x} \in \mathbb{R}_0^n$

$$V(\mathbf{x}) > 0, \quad \dot{V}(\mathbf{x}) = \sum_{i=1}^n \frac{\partial V}{\partial x_i} \frac{dx_i}{dt} = \left(\frac{\partial V}{\partial \mathbf{x}} \right)^T f(\mathbf{x}) < 0.$$

Furthermore, we also require $V(\mathbf{0}) = 0$ and that V is radially unbounded, i.e., $V(\mathbf{x}) \xrightarrow{\|\mathbf{x}\| \rightarrow \infty} \infty$. By choosing a polynomial function V that vanishes at the origin (which is a linear constraint), those last two conditions are automatically satisfied. But once again, the polynomial has to be positive, with a negative derivative. More information about the construction of such a Lyapunov function can be found in [AP15; Yan23].

Lyapunov theory also allows to design controllers for possibly unstable systems. The dynamical equations become $\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t))$, where $\mathbf{u}(t) \in \mathbb{R}^m$ is the input at time t (which could also depend on the current state). Often, one would like to minimize some cost function of the input (e.g., minimum control effort through the integral of the norm of the input $\int_0^T \|\mathbf{u}(t)\|^2 dt$), while satisfying the Lyapunov conditions [AP15]. Such a Lyapunov function is called Control Lyapunov Function [Yan23; Wee06], as its existence guarantees the stabilizability of the system in a neighbourhood of the equilibrium point. In this case, we minimize a function on the polynomial, while satisfying sum-of-squares and linear constraints.

¹A system is *globally asymptotically stable* if it converges to its equilibrium whatever the starting point is. A local version can also be expressed, as the weaker concept of stability which only requires $\dot{V}(\mathbf{x}) \leq 0$, i.e., solutions stay nearby the equilibrium point.

1.1.3 Equilibria of polynomial games

Our last example arises in game theory, when computing the Nash equilibrium of a polynomial game [Par00]. Consider a two-player zero-sum game², with a payoff function for player 1 given by

$$P(x, y) = \sum_{i=0}^n \sum_{j=0}^m p_{ij} x^i y^j,$$

where x and y are the decisions of the two players. Both players will play according to mixed strategies μ and ν in order to maximize their payoffs. A secure strategy for player 2 is to play according to a strategy ν that minimizes the best payoff player 1 can reach³:

$$\begin{aligned} \min_{\gamma, \nu} \quad & \gamma \\ \text{s.t.} \quad & \mathbb{E}_{y \sim \nu} (P(x, y)) \leq \gamma \quad \forall x \in \mathbb{R}, \\ & \int_{-\infty}^{\infty} d\nu(y) = 1. \end{aligned}$$

If player 2 plays according to that optimal strategy ν^* , player 1 is incentivised to play the x maximizing $\mathbb{E}_{y \sim \nu} (P(x, y))$, which will be upper bounded by γ . Since the game is polynomial, the expectation over ν of $P(x, y)$ can be expressed as:

$$\mathbb{E}_{y \sim \nu} (P(x, y)) = \sum_{i=0}^n \sum_{j=0}^m \mathbb{E}_{\nu} (y^j) x^i = \sum_{i=0}^n \sum_{j=0}^m \nu_j x^i,$$

where the ν_j are the moments of the measure ν . Hence, the problem becomes a minimization problem under a nonnegativity constraint of some polynomial $(\gamma - \sum_{i=0}^n \sum_{j=0}^m \nu_j x^i)$, itself depending on parameters ν_j which should correspond to the moments of some measure.

1.2 Sum-of-squares relaxation

In all the examples presented above, we are looking for nonnegative polynomials satisfying some linear constraints. Showing the nonnegativity of polynomials is therefore of great interest. A first simple necessary condition for this is that the degree d of the polynomial has to be even (in the other case, the polynomial would go to minus infinity either for $x \rightarrow \infty$, or for $x \rightarrow -\infty$). However, it appears that in the general case, the question is NP-hard [BPT12; Par00]. This can, for example, be shown by reducing it to the problem of determining whether a symmetric matrix is *copositive*⁴, which is NP-hard [Far25].

In 1900, David Hilbert published a list of 23 problems he found interesting and unsolved at that time [Hil00]. As of now, about ten of them have been solved⁵, last of which solved in 1998. Among those, the 17th problem [Rez00] asked whether every nonnegative polynomial

²A two-player zero-sum game is a game with two players, in which the first player earns exactly what the other loses and conversely.

³The second constraint corresponds to asking that ν is a valid measure.

⁴A symmetric matrix $M \in \mathbb{R}^{n \times n}$ is *copositive* if $a^T M a \geq 0$ for every nonnegative vector $a \in \mathbb{R}_{\geq 0}^n$ (i.e., with all its entries nonnegative).

⁵The exact number is difficult to determine, as mathematicians have not reached a consensus on some of them.

can be expressed as a sum of squared rational functions. Hilbert himself gave an answer to a similar question: can every nonnegative polynomial be written as a sum of squared polynomials? Mathematically, he asked himself whether, for all nonnegative polynomial $p \in \mathbb{R}[x_1, \dots, x_n]_d$, there exists a finite number of polynomials $u_i \in \mathbb{R}[x_1, \dots, x_n]_{d/2}$, $i = 1, \dots, r$, such that

$$p = \sum_{i=1}^r u_i^2.$$

Such polynomials are called *sum of squares* (SOS), as they can be expressed as the sum of squared polynomials.

The answer to the question is yes if and only if the polynomial is either a quadratic form (i.e., a polynomial of degree 2), or a univariate polynomial (or equivalently a bivariate *homogeneous* polynomial⁶), or a bivariate quartic form (i.e., a polynomial in two variables of degree 4). A well-known counterexample for a nonnegative polynomial which cannot be written as a sum of squares is the Motzkin polynomial⁷ [BPT12], given by $x_1^4 x_2^2 + x_1^2 x_2^4 - 3x_1^2 x_2^2 + 1$.

Why then are sum-of-squares polynomials of interest? The answer is simple: every sum-of-squares polynomial is nonnegative. Furthermore, determining whether a polynomial is a sum of squares or not can be expressed as a semidefinite program [BPT12; Far25], which can be solved in polynomial time. For this reason, people often consider relaxations of the different problems presented above by asking the polynomials to be sum of squares instead of nonnegative. Furthermore, for feasibility problems like the existence of a Lyapunov function (Section 1.1.2), a sum-of-squares polynomial gives a certificate for the controllability of a system, even if it is not optimal.

1.3 Main problem

This motivates us to analyse a linear sum-of-squares optimization problem of the kind

$$\begin{aligned} \min_{p \in \mathbb{R}[\mathbf{x}]_{2d}} \quad & \mathcal{L}_1(p(\mathbf{x})) \\ \text{s.t.} \quad & \mathcal{L}_2(p(\mathbf{x})) = \mathbf{b}, \\ & p \text{ is a sum of squares,} \end{aligned} \tag{1.1}$$

for some **surjective** linear operators $\mathcal{L}_1 : \mathbb{R}[\mathbf{x}]_{2d} \rightarrow \mathbb{R}$ and $\mathcal{L}_2 : \mathbb{R}[\mathbf{x}]_{2d} \rightarrow \mathbb{R}^m$, and a constraint vector $\mathbf{b} \in \mathbb{R}^m$. More than asking for a polynomial to be SOS, we have a linear objective and linear constraints on the polynomial. As will be detailed in Section 2.3, this problem can be modelled as a semidefinite program (SDP), and solved with classical optimization algorithms such as the Interior Point Method [Nes18; Nes00].

However, the number of variables of the SDP (1.1) grows quadratically with the dimension of the set of multivariate polynomials in n variables with degree d , itself growing as $\binom{n+d}{d}$ [Yan24]. For a bivariate polynomial of degree 8, we are already at more than

⁶*homogeneous* means that all the terms are of the same degree.

⁷This example is also known as the first example of a nonnegative polynomial that cannot be written as a sum of squares.

4000 variables. The issue of the high number of variables is common to all SDPs, making them time and memory intractable for many problems [BM03].

To address this issue, Burer and Monteiro proposed a nonlinear algorithm for solving SDPs [BM03]⁸ at the turn of the 21st century. Their approach consists of factorizing the variable matrix $X \in \mathbb{S}_+^n$, the set of positive semidefinite (PSD) matrices of size n , into a product VV^T , where $V \in \mathbb{R}^{n \times r}$. One motivation behind this choice is inspired by a result of Pataki [Pat98] and Barvinok [Bar95], which showed that for every SDP with m constraints and a compact search space, there exists an optimal solution of rank r , with r satisfying $r(r+1) \leq 2m$. This suggests the existence of a low-rank solution to SDPs, which can be represented mathematically by choosing a small number of columns r for V .

Applying the idea of Burer and Monteiro on SOS programs like (1.1) corresponds to expressing $p(\mathbf{x})$ as a sum of polynomials $q_1(\mathbf{x}), \dots, q_r(\mathbf{x})$ squared. For univariate polynomials, we even have the result that there exists a rank-2 solution, i.e., with two polynomials ($r=2$) [CLR94]! This, as well as other recent results on univariate polynomials [LYP23], motivated us to first consider univariate polynomials. A generalization to multivariate polynomials will then be considered (not in this thesis); but as shown in [Ble+24], working with multivariate polynomials is not that simple, requiring advanced notions of algebraic geometry.

Inspired by Burer and Monteiro, we reformulate (1.1) for univariate polynomials as

$$\begin{aligned} \min_{(u_1, u_2) \in \mathbb{R}[x]_d^2} \quad & \mathcal{L}_1(u_1(x)^2 + u_2(x)^2) \\ \text{s.t.} \quad & \mathcal{L}_2(u_1(x)^2 + u_2(x)^2) = \mathbf{b}. \end{aligned} \tag{1.2}$$

The notation $\mathbb{R}[x]_d^2$ corresponds to the set of pairs of degree- d polynomials. It is important to note that the decision variables are the polynomials $u_1, u_2 \in \mathbb{R}[x]_d$, but the objective and constraints are linear operators on the SOS polynomial $u_1^2 + u_2^2$. Those are therefore quadratic, and potentially nonconvex. Problem (1.2) will be our common thread over the manuscript, and will be called main problem. The word *univariate* will often be omitted in the following chapters.

To illustrate the formalism of (1.2), let us return to our motivating example of finding the best lower bound of a polynomial (Section 1.1.1). In that case, the operator $\mathcal{L}_1$ will output the constant term of $u_1^2(x) + u_2^2(x)$, and the constraints will look like⁹ $\frac{d}{dx}(u_1^2(x) + u_2^2(x)) = p'(x)$ (which is a set of $2d$ linear constraints, the polynomial $p \in \mathbb{R}[x]_{2d}$ being fixed).

1.4 A Riemannian approach

In their paper, Burer and Monteiro [BM03; BM05] proposed an Augmented Lagrangian Method to solve the nonlinear problem resulting from their reformulation. Augmented Lagrangian approaches have been widely studied in the last decades, as they allow one to

⁸This paper earned the *SIAM Activity Group on Optimization Test of Time Award* in 2023, to award the significant and sustained influence of the paper over a long time period [Ano].

⁹Only the constant term between the two sides of the equality may differ, due to the introduction of the variable γ .

tackle any constrained optimization problem, without caring too much about the structure of the feasible set [BV04; NW06; WM22]. Such methods generally consist in reformulating the problem as an unconstrained one, but with a penalty on constraint satisfaction and a so-called Lagrangian term, both aiming to force the satisfaction of the constraints. The new problem is then solved by iteratively minimizing over the primal variables, with an update of the dual variables, corresponding to a penalty parameter and Lagrangian multipliers following each iteration. The main advantage of such methods is that they are straightforward to apply on new problems.

While Burer and Monteiro demonstrated their Augmented Lagrangian algorithm performs well in comparison to classical SDP approaches, it does not exploit the intrinsic properties of the problems. Yet, one can expect much better performances for an algorithm that would be specialized for the kind of problem we consider (in particular (1.2) in our case). This is the case of Riemannian Optimization Algorithms [AMS08; Bou23]. Those algorithms are designed to optimize on nonlinear smooth sets, called manifolds, by finding iterates lying on the manifold. This is made possible by the use of Riemannian gradients, a generalization of the concept of gradient, and retractions, ensuring that all the iterates lie on the manifold. Both of those concepts are specific to the manifold on which the optimization process takes place, and therefore require a thorough understanding of the feasible set.

In addition to the advantage of more specialized algorithms, understanding the feasible set enables us to uncover properties of the problem. In recent years, several researchers have looked into the issue of benign nonconvexity, i.e., nonconvex problems without local non-global extrema. This is the case of N. Boumal et al. [BVB18], who showed that the Burer and Monteiro approach for SDP programs has no spurious local minima (i.e., local minima that are not global) when the feasible set of the nonlinear reformulation is a smooth embedded submanifold with linear independence of the gradients of the constraints at each point¹⁰. In 2022, Legat, Yuan and Parrilo showed in [LYP23] that the problem of finding a sum-of-squares decomposition of a polynomial through a quadratic penalization of the constraint (i.e. $\min_{\mathbf{u} \in \mathbb{R}[x]_d^2} \|\sigma(\mathbf{u}) - p\|$, where p is the given polynomial) has no spurious local minima neither. This allows them to find a feasible point to the problem (1.2) in near linear time. The goal of this Thesis is to deepen the understanding of the feasible set already developed in their paper, in order to use techniques from Riemannian optimization to solve the problem (1.2).

Moreover, it is worth noting that a widely used technique to solve sum-of-squares optimization problems of the kind (1.1) is through their dual [Las10; BPT12]. The dual cone of the space of nonnegative polynomials is the set of linear functionals that take nonnegative values for all nonnegative polynomials. A nice interpretation of this dual cone is in terms of truncated moment sequences of probability distributions, as we can identify the dual space with the set of moments for which a nonnegative measure matching those moments exists. On the dual formulation, Lasserre et al. [Las01; HL05] developed an algorithm solving a sequence of Linear Matrix Inequality (LMI) problems of increasingly larger size to determine whether the truncated sequence of moments corresponds to a nonnegative measure. While those LMIs are only necessary conditions, they also developed a sufficient condition for global optimality, called the *flatness condition*, allowing the algorithm to stop after a few iteration in many cases.

¹⁰They also require the feasible set of the original SDP formulation to be compact.

Finally, let us also mention the existence of the Projected Gradient Method and Sequential Quadratic Programming for solving problem (1.1). The first one requires the description of a projection on the set of nonnegative polynomials, and therefore a good understanding of the feasible set too. The second one solves sequentially quadratic programming subproblems in order to converge to a critical point [NW06, Chapter 18]. While it would be interesting to analyze those methods, it has not been done in this Thesis.

1.5 Structure of the work

In the [second chapter](#), we recall some notions on nonnegative and sum-of-squares polynomials. To the best of our knowledge, a manifold structure for the set of squared polynomials has never been analyzed. Accordingly, we establish a topological foundation for the set of univariate polynomials in Chapter 3. This allows us to show that the pairs of coprime univariate polynomials form an open submanifold of the pairs of univariate polynomials. In the [fourth chapter](#), we explore in more detail the feasible set of the problem (1.2), by describing a few cases in which we can conclude it is a smooth embedded submanifold. However, this is not always the case. In Chapter 5, we present an example of a problem (1.2) whose feasible set is not a smooth manifold, but a stratified set. We also obtain an expression of the tangent cone to this feasible set. Finally, in Chapter 6, we describe a first-order optimization algorithm for solving linear sum-of-squares optimization problems for which the feasible set is a smooth embedded submanifold. In particular, we show that the Newton method for finding the zeros of a function can be used as a retraction for embedded submanifolds of Euclidean spaces.

1.6 Contribution

- In Chapter 3, we establish a topology on the pairs of polynomials of degree d , $\mathbb{R}[x]_d^2$. Next, we show that its restriction to the pairs of coprime polynomials is an open submanifold, relying on the resultant between two polynomials;
- In Chapter 4, we show that when the linear operator $\mathcal{L}_2$ defining the feasible set of (1.2) is bijective, the feasible set is a smooth manifold. Moreover, the restriction of the feasible set of (1.2) to the set of coprime pairs (potentially multiplied by some polynomial) is also a smooth embedded submanifold;
- In Chapter 5, we present an example of problem of the kind (1.2) for which the feasible set is not a smooth manifold. This is proved by relying on the expression of the tangent cone we develop in Section 5.2. Yet, on this example, the feasible set is shown to be a stratified set;
- In Chapter 6, we introduce a Riemannian Gradient Descent algorithm on smooth embedded submanifolds of $\mathbb{R}[x]_d^2$, explaining how to obtain an expression of the Riemannian gradient. Furthermore, we show that the Newton-Raphson method to get a zero of a local defining function is a smooth retraction. This may be of independent interest as it is valid on any smooth embedded submanifold (of a Euclidean space) that can be locally described by a local defining function.

Chapter 2

Nonnegative and Sum-of-Squares polynomials

In this chapter, we review some theoretical background on nonnegative and sum-of-squares polynomials¹. In particular, we show the equivalence between nonnegativity and sum-of-squares (Lemma 2.2) for univariate polynomials. In Section 2.2, we review the concept of basis for univariate polynomials. Next, in Section 2.3, we present the classical SDP formulation for the problem (1.1). At the end of the chapter, we explain how to find efficiently a feasible solution to our main problem (1.2). For this, we rely on a recent paper of Legat et al. [LYP23], in which they show how to obtain such a point with a first-order optimization algorithm such as L-BFGS.

Interesting references on polynomials and sum-of-squares include [Yan24; BPT12; CLO25].

2.1 Nonnegative and sum-of-squares polynomials

As already mentioned in Section 1.3, every sum-of-squares polynomial is nonnegative. Yet, the converse is not true in general. Hilbert showed in the late 19th century the equivalence between nonnegative and sum-of-squares polynomials holds only in three specific cases: univariate polynomials (see next lemma); quadratic polynomials (i.e., of degree 2, whatever the number of variables is); and quartic bivariate polynomials (i.e., degree-4 polynomials in two variables) [Rez00]. However, even if the non-equivalence between those two sets in the other cases had been fully established, no example of nonnegative polynomial that is not SOS had been found until the mathematician T. Motzkin proposed the following, in the mid-20th century: $x_1^4 x_2^2 + x_1^2 x_2^4 - 3x_1^2 x_2^2 + 1$. A proof of its nonnegativity and that it is not SOS is given in [BPT12]. Since the result for univariate polynomials will be of major importance in Chapter 4, we restate it in Lemma 2.2 with a proof.

But before going into the proof of that lemma, let us first recall a few fundamental concepts about polynomials. The fundamental theorem of algebra states that every

¹Concerning the theoretical notions on Riemannian optimization, we have decided to introduce them when necessary in the following sections.

univariate polynomial $p \in \mathbb{R}[x]_d$ (of degree exactly d) can be factorized as

$$p(x) = \sum_{i=0}^d c_i x^i = c_d \prod_{i=1}^d (x - x_i),$$

where $x_i \in \mathbb{C}$ are the roots of the polynomial. This factorization is unique up to reordering of the factors. Given two polynomials $p_1 \in \mathbb{R}[x]_{d_1}$ and $p_2 \in \mathbb{R}[x]_{d_2}$, we define their *greatest common divisor* $g = \gcd(p_1, p_2)$ as the highest degree polynomial that divides both of them. By *divides*, we mean that there exist polynomials w_1, w_2 such that $p_1 = g \cdot w_1$ and $p_2 = g \cdot w_2$. The greatest common divisor of two polynomials is unique up to multiplication by a scalar. We thus fix it to the polynomial with a unitary leading coefficient². Last but not least, two polynomials are said to be *coprime* if their greatest common divisor is 1, i.e., if they do not share a common root.

Remark 2.1. It may be confusing to write $p(x)$ to mean the polynomial p in the variable x or the polynomial p evaluated at the point x . However, the distinction should be clear from the context. $\triangle$

Lemma 2.2 – Nonnegative univariate polynomials are SOS

Let $p(x) \in \mathbb{R}[x]_{2d}$ be a univariate polynomial.

- (a) The polynomial $p(x)$ is nonnegative if and only if it is a sum of (two) squared polynomials in $\mathbb{R}[x]_d$;
- (b) Furthermore, if the number of pairs of complex roots of $p(x)$ is $K \geq 1$, there are at most 2^{K-1} different ways of writing $p(x)$ as the sum of two squares, up to multiplication by a rotation matrix.

Proof. This proof extends the one given in [CLR94] by explicitly describing the set of pairs of polynomials whose sum of squares equals $p(x)$ as a union of “circles”.

(a) The direction $\text{SOS} \Rightarrow \text{nonnegative}$ is straightforward.

Assume $p(x)$ is a nonnegative polynomial. By the fundamental theorem of algebra, we can factorize $p(x)$ as

$$p(x) = a \cdot \prod_{i=1}^R (x - r_i) \cdot \prod_{j=1}^C (x - c_j) =: a \cdot r(x) \cdot c(x),$$

where $r_i \in \mathbb{R}$ (resp. $c_j \in \mathbb{C}$) are the real (resp. complex) roots of $p(x)$, $R + C = 2d$, and $a \in \mathbb{R}$.

We note that each real root r_i must be of even multiplicity. By contradiction, if it has an odd multiplicity, then there exists a neighbourhood of r_i with no other root, and on which the polynomial changes sign, contradicting the fact that it is nonnegative. Therefore, $r(x)$ is a squared polynomial.

Next, each complex root c_j has to come with its complex conjugate since $p(x)$ is a polynomial with real coefficients: $C = 2K$. Therefore, one can write (after reordering of

²The polynomial is said to be *monic* if its leading coefficient is one.

the terms):

$$\begin{aligned} c(x) &= \prod_{j=1}^K (x - c_j) \cdot \prod_{j=1}^K (x - \bar{c}_j) = \prod_{j=1}^K (x - c_j) \cdot \overline{\prod_{j=1}^K (x - c_j)} \\ &=: (u_1(x) + iu_2(x)) \cdot \overline{(u_1(x) + iu_2(x))} = u_1^2(x) + u_2^2(x). \end{aligned} \quad (2.1)$$

This shows that $c(x)$ can be expressed as a sum of two squares.

Since $p(x)$, $r(x)$ and $c(x)$ are nonnegative polynomials, the leading coefficient a has to be nonnegative (its square root is therefore well defined). This allows to conclude that $p(x)$ is a sum of (two) squares, being the product of a nonnegative real a , a squared polynomial $r(x)$, and a sum of two squared polynomials $c(x)$.

(b) Now, we would like to describe the set of decompositions of $p(x)$ as the sum of two squares (i.e., $u^2(x) + v^2(x)$ for some polynomials $u(x), v(x) \in \mathbb{R}[x]_d$). At the real roots r_i of $p(x)$, both polynomials of the decomposition must vanish³. Hence, $u(x)$ and $v(x)$ are multiples of $r(x)$, and the only difference between the SOS decompositions can appear in the decomposition of $c(x)$ as the sum of two squares. Let us swap c_j (for some j) with its complex conjugate $\bar{c}_j$ in (2.1). This gives us another decomposition of $c(x)$ as the sum of two polynomials $u_1'(x)$ and $u_2'(x)$ squared. By considering all the configurations of the roots in two groups (like in (2.1)), the maximal number of such sum of squares decompositions is 2^{K-1} . This number is reached when all the complex conjugate pairs of roots are distinct. Note that symmetric configurations have been removed, as considering a combination of roots or the combination of the complex conjugates of all those roots yields the same SOS decomposition (which explains the exponent $K - 1$ instead of K).

It remains to show that, up to multiplication by a rotation matrix, every SOS decomposition of $c(x)$ is one of the (at most) 2^{K-1} previously found. Assume we have a decomposition of $c(x)$ as the sum of the squares of $\tilde{u}_1(x)$ and $\tilde{u}_2(x)$:

$$c(x) = \tilde{u}_1(x)^2 + \tilde{u}_2(x)^2 = (\tilde{u}_1(x) + i\tilde{u}_2(x))(\tilde{u}_1(x) - i\tilde{u}_2(x))$$

Since a univariate polynomial can be factorized in one unique form, we must have⁴

$$(\tilde{u}_1(x) + i\tilde{u}_2(x)) = \lambda(u_1'(x) + iu_2'(x))$$

for some $(u_1'(x), u_2'(x))$ obtained previously and $\lambda \in \mathbb{C}$ such that⁵ $\lambda\bar{\lambda} = 1$. Therefore, λ lies on the unit complex circle, and $(\tilde{u}_1(x), \tilde{u}_2(x))$ is a rotation of $(u_1'(x), u_2'(x))$: there exists $\theta \in [0, 2\pi[$ such that:

$$\begin{pmatrix} \tilde{u}_1(x) \\ \tilde{u}_2(x) \end{pmatrix} = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} u_1'(x) \\ u_2'(x) \end{pmatrix}.$$

□

Remark 2.3. The last step of the proof also works with a reflection matrix. However, this simply corresponds to switching $u_1'(x)$ and $u_2'(x)$ before applying a rotation. △

³If some roots have multiplicity larger than one, then the derivatives must also vanish.

⁴Indeed, all the permutations of the roots have been explored in (2.1). However, it could be equal to $\lambda(u_1'(x) - iu_2'(x))$, but this case is similar, and corresponds to the multiplication by a reflection matrix (see Remark 2.3).

⁵Since $c(x) = (\tilde{u}_1(x) + i\tilde{u}_2(x))(\tilde{u}_1(x) - i\tilde{u}_2(x)) = \lambda\bar{\lambda}(u_1'(x) + iu_2'(x))(u_1'(x) - iu_2'(x)) = \lambda\bar{\lambda}c(x)$.

Remark 2.4. An easier way to see that a rotation of an SOS decomposition of a polynomial is still an SOS decomposition of that polynomial is the following: consider a nonnegative polynomial $p(x) = u_1^2(x) + u_2^2(x)$; then, by the Pythagorean identity, multiplying (u_1, u_2) by a rotation matrix gives again an SOS decomposition of p :

$$(\cos(\theta)u_1 - \sin(\theta)u_2)^2 + (\sin(\theta)u_1 + \cos(\theta)u_2)^2 = u_1^2 + u_2^2.$$

△

Remark 2.5. Lemma 2.2 states that, up to multiplication by a rotation matrix, a nonnegative univariate polynomial p can be expressed as the sum of two squared polynomials in at most 2^{K-1} different ways, where K is the number of pairs of complex conjugate roots of p . The rotation matrix allows for an analogy with circles that will be used several times throughout this thesis. △

2.2 Basis for univariate polynomials

For linear spaces, a basis is a set of vectors such that every element of the linear space can be written in a unique way as a linear combination of those basis vectors. Furthermore, two distinct elements will have different representations in the basis. This is interesting as it means that two polynomials are equal if and only if their expression in a chosen basis is the same.

Several bases have been studied in the literature for the vector space of univariate polynomials of degree at most d , $\mathbb{R}[x]_d$, which is $d + 1$ dimensional [Emb16]. The simplest basis one can consider is the monomial basis, given by all the monomials in $\mathbb{R}[x]_d$: $\{1, x, \dots, x^d\}$. Another interesting basis is the Lagrangian one. Given a set of $d + 1$ distinct points $\{x_k\}_{k=0}^d \subset \mathbb{R}$, we define the Lagrangian basis $\{l_k\}_{k=0}^d$ as:

$$l_k(x) = \prod_{\substack{0 \leq i \leq d \\ i \neq k}} \frac{x - x_i}{x_k - x_i}, \quad k = 0, \dots, d.$$

A key property of the Lagrangian basis is that every polynomial $l_k(x)$ is zero at $x_0, \dots, x_{k-1}, x_{k+1}, \dots, x_d$, and has value one at x_k . Therefore, every polynomial $p \in \mathbb{R}[x]_d$ is fully described by its values at $\{x_k\}_{k=0}^d$:

$$p(x) = \sum_{k=0}^d p(x_k) l_k(x).$$

Note that the basis depends on the choice of the $\{x_k\}_{k=0}^d$. Given a polynomial, it can be interesting to consider the Lagrangian polynomials defined at its roots (and other points to complete the basis). Indeed, expressing the polynomial in that basis will result into a sparse vector. This property has already been used in the proof of Lemma 2.2: if two polynomials are equal, their expression in the Lagrangian basis is identical, which means that they have the same value at the $d + 1$ chosen points. Hence, if a polynomial vanishes at some point, all polynomials that are equal to that polynomial must vanish at that point. This was used when we said that $u(x)$ and $v(x)$ must vanish at the real roots r_i : $0 = p(r_i) = u(r_i)^2 + v(r_i)^2$. Since the two terms on the right-hand side are nonnegative, they must be zero.

While the monomial basis will be of interest when constructing a topology on the set of pairs of polynomials (Chapter 3), the Lagrangian one will be used for describing linear applications in Chapter 5.

2.3 Linear Sum-of-Squares Programming

The standard approach for solving a linear SOS program of the kind (1.1) is to reformulate it as a semidefinite program (SDP). Let β be a vector of polynomials forming a basis for $\mathbb{R}[x]_d$ (e.g., the monomial or Lagrangian described above). Given a polynomial $p \in \mathbb{R}[x]_{2d}$, it is a sum of squares if and only if there exist polynomials $u_1, \dots, u_r \in \mathbb{R}[x]_d$ corresponding to an SOS decomposition⁶ of p :

$$p = \sum_{i=1}^r u_i^2.$$

By denoting $\mathbf{u} \in \mathbb{R}[x]_d^r$ the vector of polynomials u_i , there exists a (unique) matrix $V \in \mathbb{R}^{r \times (d+1)}$ such that $\mathbf{u} = V\beta$. The condition for being SOS can therefore be written as

$$p = (V\beta)^T V\beta = \beta^T Q\beta,$$

for some positive semidefinite (PSD) matrix $Q = V^T V \in \mathbb{S}_+^{d+1}$, the set of $(d+1) \times (d+1)$ PSD matrices. The rank of Q corresponds to the number of polynomials in the decomposition.

Lemma 2.6 – Sum-of-squares as SDP [Yan24]

The linear sum-of-squares program (1.1) can be formulated as an SDP.

Proof. We have seen that a polynomial $p \in \mathbb{R}[x]_{2d}$ is SOS if and only if there exists a positive semidefinite matrix $Q \in \mathbb{S}_+^{d+1}$ such that $p = \beta^T Q\beta$, for some basis β of $\mathbb{R}[x]_d$. Those constraints are in fact linear on the entries of Q , and may be interpreted as matching the coefficients of the polynomials on the two sides of the equality. There exists therefore a linear map $\mathcal{L} : \mathbb{S}_+^{d+1} \rightarrow \mathbb{R}[x]_{2d} : Q \mapsto p(x) = \beta^T Q\beta$, mapping an SDP matrix Q into its associated nonnegative polynomial. The composition of the linear operators $\mathcal{L}_1$ and $\mathcal{L}_2$ with $\mathcal{L}$ gives in turn a linear application on the entries of Q . For completeness, we write the equivalent program of (1.1) as SDP an:

$$\begin{aligned} \min_{Q \in \mathbb{S}_+^{d+1}} & (\mathcal{L}_1 \circ \mathcal{L})(Q) \\ \text{s.t.} & (\mathcal{L}_2 \circ \mathcal{L})(Q) = \mathbf{b}. \end{aligned}$$

□

2.4 Nonlinear program for the feasibility problem

As already mentioned in the introduction, SDPs quickly become intractable in terms of both time and memory, due to the number of variables that grows quadratically with the degree d of the polynomials (in the univariate case). Burer and Monteiro proposed an

⁶This decomposition may be not unique as shown in Lemma 2.2

approach to solve SDP problems, replacing the unknown matrix Q by $V^T V$, where V is a rectangular matrix whose number of rows is equal to the rank of Q [BM03; BM05]. While this reformulation introduces nonlinearity and generally nonconvexity, it is interesting as it drastically reduces the number of variables.

For the univariate sum-of-squares problem, we showed in Lemma 2.2 that there exists a rank-2 solution (i.e., with only two polynomials in the decomposition). Legat et al. exploited that idea in [LYP23] by analyzing the problem of finding an SOS decomposition of a nonnegative polynomial $p \in \mathbb{R}[x]_{2d}$ through the minimization of a quadratic penalty:

$$\min_{\mathbf{u} \in \mathbb{R}[x]_d^r} f_p(\mathbf{u}) := \left\| \sum_{i=1}^r u_i^2 - p \right\|^2.$$

The authors show (Theorem 1.1) that the function $f_p(\mathbf{u})$ has no *spurious local minima* for $r = 2$, i.e.; all the local minima are global⁷. Hence, finding an SOS decomposition of a polynomial $p \in \mathbb{R}[x]_{2d}$ can be done efficiently by minimizing the function f_p with a first-order algorithm, like L-BFGS [NW06, Section 7.2]. Using an FFT algorithm to compute the gradient of f_p , their results show a linear convergence and a near linear time complexity.

In Section 6.3 of the paper, they show that this result generalizes to the problem of finding a feasible point $p \in \mathbb{R}[x]_{2d}$ in the intersection of the cone of SOS polynomials (i.e., p should be SOS) with an affine subspace: $\mathcal{L}(p) = \mathbf{b}$. This corresponds to the feasible set of our problem (1.1). The quadratic penalty function becomes $f(\mathbf{u}) := \left\| \mathcal{L} \left(\sum_{i=1}^2 u_i^2 \right) - \mathbf{b} \right\|^2$. In Chapter 6, we will design a *first-order Riemannian optimization* algorithm on *embedded submanifolds* (see Chapters 3 and 4) of $\mathbb{R}[x]_d^2$, the set of pairs of degree- d polynomials. To apply that algorithm, it will be necessary to get a first iterate on the manifold, motivating the use of a minimization algorithm on the quadratic penalty f .

⁷Yuan generalized the results for all $r \geq 2$ in his Ph.D. Thesis [Yua22].

Chapter 3

Coprime polynomials as an open submanifold

Having covered the basics of polynomials in the [previous chapter](#), we may now define a manifold structure on polynomials. A smooth manifold can be seen as a “smooth” surface, on which calculus theory, like smooth maps, differentials, gradients, etc. can be generalized. N. Boumal defines it in broad terms as a set that, locally, “can be smoothly deformed to a linear subspace, and back” [\[Bou26\]](#). Intuitively, a sphere or a torus are smooth manifolds, but a cross is not.

Relying on a range of references in Riemannian optimization [\[AMS08; Lee12; Bou23\]](#), we demonstrate that the set of pairs of coprime univariate degree- d polynomials, $\mathbb{R}[x]_d^{2,c}$, forms an *open submanifold* of the set of pairs of univariate degree- d polynomials, $\mathbb{R}[x]_d^2$. For the concept of openness to be well defined, we first establish in Section [3.1](#) a natural topology on $\mathbb{R}[x]_d^2$ by mapping polynomials to vectors in $\mathbb{R}^n$ via their coefficients. Following that, in Section [3.2](#), we use the notion of *resultant* between polynomials to show that $\mathbb{R}[x]_d^{2,c}$ is an open subset of $\mathbb{R}[x]_d^2$, making it an open submanifold.

In the [next chapter](#), we will define more complex *embedded submanifolds* of $\mathbb{R}[x]_d^2$. Those will correspond to the feasible sets of our main problem [\(1.2\)](#), for specific linear operators $\mathcal{L}_2$ and constraint vectors $\mathbf{b}$ defining the feasible sets. Coprimality between polynomials will be of great interest, as it removes limitations that non-coprime polynomials may have.

3.1 Topology on pairs of polynomials

As a starting point, and without entering into the detail, we recall that a vector space is a smooth manifold¹. As previously explained, a smooth manifold can be locally deformed to a linear subspace $\mathbb{R}^k$. The *dimension* of a manifold is defined as the dimension k of that linear subspace. For a vector space, this corresponds to the dimension of the vector space. In particular, $\mathbb{R}[x]_d^2$, the set of pairs of degree at most d polynomials, is a smooth manifold of dimension $2(d+1)$.

We define an *open submanifold* $\mathcal{N}$ of a manifold $\mathcal{M}$ as an open subset of it:

¹Intuitively, this makes sense as smooth manifolds are surfaces on which calculus theory can be generalized, but calculus theory is already well-defined on vector spaces through an isomorphism with linear spaces.

Definition 3.1 – Open submanifold

Let $\mathcal{M}$ be a smooth manifold. Every (non-empty) open subset $\mathcal{N}$ of $\mathcal{M}$ is called an *open submanifold* of $\mathcal{M}$, and has the same dimension as $\mathcal{M}$.

In Definition 3.1, we use the notion of ‘open’ subset of $\mathcal{M}$, which is inherently paired with the concept of *topology* [AMS08, Appendix A.2].

Definition 3.2 – Topology

A *topology* on a set X is a collection $\mathcal{T}$ of subsets of X , called *open sets*, such that:

1. X and $\emptyset$ belong to $\mathcal{T}$;
2. the union of the elements of every subcollection of $\mathcal{T}$ is in $\mathcal{T}$;
3. the intersection of the elements of every finite subcollection of $\mathcal{T}$ is in $\mathcal{T}$.

Like every finite-dimensional vector space over the reals, $\mathbb{R}[x]_d$ inherits a natural topology via an isomorphism $F : \mathbb{R}[x]_d \rightarrow \mathbb{R}^{d+1}$ (the choice of $\mathbb{R}^{d+1}$ follows from the fact that $\mathbb{R}[x]_d$ is $(d+1)$ -dimensional), by defining a subset X of $\mathbb{R}[x]_d$ to be open if and only if $F(X)$ is open in $\mathbb{R}^{d+1}$. For this, we may use the representation of a vector in a basis of $\mathbb{R}[x]_d$, which is an isomorphism. Indeed, by definition of basis, every element of $\mathbb{R}[x]_d$ can be expressed in a unique way as a linear combination of the elements in the basis.

Among the bases defined in Section 2.2, the monomial basis will be particularly suitable to show that $\mathbb{R}[x]_d^{2,c}$, the set of pairs of coprime polynomials of degree at most d , is an open submanifold of $\mathbb{R}[x]_d^2$. The choice of the monomial basis corresponds to mapping an element of $\mathbb{R}[x]_d$ into the vector of its coefficients. For $\mathbb{R}[x]_d^2$, we may simply stack the representations of the two polynomials in a vector of size $2(d+1)$, and define a similar isomorphism.

3.2 Manifold of coprime polynomials

Now that the topology on the pairs of degree- d polynomials has been defined, we may proceed with showing that $\mathbb{R}[x]_d^{2,c}$ is open in $\mathbb{R}[x]_d^2$. For this, we rely on the notion of *resultant* of two polynomials [Woo16], which is defined below.

Definition 3.3 – Sylvester matrix

Given two univariate polynomials $p(x) = \sum_{i=0}^m a_i x^i$ and $q(x) = \sum_{i=0}^n b_i x^i$, with a_m and b_n nonzero, we define the $(n+m) \times (n+m)$ *Sylvester matrix*

$$\text{Syl}(p, q) := \begin{pmatrix} a_m & & & & b_n & & & \\ a_{m-1} & a_m & & & b_{n-1} & b_n & & \\ \vdots & \vdots & \ddots & & \vdots & \vdots & \ddots & \\ a_1 & a_2 & \cdots & a_m & b_2 & b_3 & \cdots & \\ a_0 & a_1 & \cdots & a_{m-1} & b_1 & b_2 & \cdots & b_n \\ & a_0 & \cdots & a_{m-2} & b_0 & b_1 & \cdots & b_{n-1} \\ & & \ddots & \vdots & & \ddots & \ddots & \vdots \\ & & & a_0 & & & & b_0 \end{pmatrix},$$

whose first n (resp. last m) columns contain coefficients of p (resp. q).

Definition 3.4 – Resultant

Given two univariate polynomials $p(x) = \sum_{i=0}^m a_i x^i$ and $q(x) = \sum_{i=0}^n b_i x^i$, with a_m and b_n nonzero, their *resultant* is defined as the determinant of the corresponding Sylvester matrix:

$$\text{Res}(p, q) = |\text{Syl}(p, q)|. \quad (3.1)$$

One of the main properties of the resultant is its connection with the roots of the two polynomials, and specifically with the differences between those roots.

Theorem 3.5

Given two univariate polynomials $p(x) = \sum_{i=0}^m a_i x^i$ and $q(x) = \sum_{i=0}^n b_i x^i$, with a_m and b_n nonzero, their resultant is proportional to the product of the difference of their roots:

$$\text{Res}(p, q) = a_m^n b_n^m \prod_{i,j} (\alpha_i - \beta_j), \quad (3.2)$$

where α_i and β_j are the roots of p and q respectively.

Proof. See [CLO05, Section 3.1]. □

Corollary 3.6. *Two univariate polynomials share a common root if and only if their resultant is zero. Conversely, two univariate polynomials are coprime if and only if their resultant is nonzero.*

Proof. The necessary (only if) condition follows directly from the previous theorem. Indeed, by contradiction, if two univariate polynomials are not coprime, (3.2) indicates that their resultant is zero.

The sufficient condition is slightly more involved, and relies on the fact that the roots of polynomials with coefficients in $\mathbb{R}$ belong to the field $\mathbb{C}$. As a field is an integral domain [CLO25, Def A.3], it has no zero divisor, meaning that a product of factors in the field is equal to zero if and only if at least one of those factors is zero:

$$a \cdot b = 0 \Leftrightarrow a = 0 \text{ or } b = 0.$$

Therefore, the leading coefficients being nonzero, at least one of the roots of p must be equal to a root of q for (3.2) to be zero. $\square$

This dual definition on the resultant, which on the one hand expresses the resultant as the determinant of the Sylvester matrix (3.1) (a polynomial in the coefficients of the two polynomials), and on the other as the product of the differences of their roots (3.2), give a concise proof of the openness of $\mathbb{R}[x]_d^{2,c}$ in $\mathbb{R}[x]_d^2$.

Theorem 3.7

The set of pairs of coprime univariate degree- d polynomials $\mathbb{R}[x]_d^{2,c}$ is an open subset of $\mathbb{R}[x]_d^2$.

Proof. The set of non-coprime pairs of polynomials consists of the pairs of polynomials whose resultant equals zero. Since the resultant of two polynomials is a polynomial expression of their coefficients, the preimage of the closed set $\{0\}$ under the resultant map is a closed set in $\mathbb{R}[x]_d^2$. Therefore, its complement, which corresponds to $\mathbb{R}[x]_d^{2,c}$, is open. $\square$

It follows by Definition 3.1 that $\mathbb{R}[x]_d^{2,c}$ is an open submanifold of $\mathbb{R}[x]_d^2$.

Remark 3.8. The closure of $\mathbb{R}[x]_d^{2,c}$ is all $\mathbb{R}[x]_d^2$. Indeed, given two non-coprime polynomials, a small perturbation of one of them modifies its roots, so that they no longer coincide with those of the other polynomial, making the resultant nonzero. Therefore, all the pairs of non-coprime polynomials belong to the closure of $\mathbb{R}[x]_d^{2,c}$. $\triangle$

Chapter 4

Manifold structures on the feasible set of sum-of-squares problem

With the algebraic and topological groundwork in place, we are now ready to study the feasible set of our main problem (1.2) in more detail. Remember, we are looking for pairs of polynomials in $\mathbb{R}[x]_d$ such that their sum of squares satisfies a linear set of equations:

$$\mathcal{S} := \left\{ (u_1, u_2) \in \mathbb{R}[x]_d^2 \mid \mathcal{L}_2(u_1^2(x) + u_2^2(x)) = \mathbf{b} \right\}, \quad (4.1)$$

where $\mathcal{L}_2 : \mathbb{R}[x]_{2d} \rightarrow \mathbb{R}^m$ is a surjective linear operator and $\mathbf{b} \in \mathbb{R}^m$ is a constraint vector.

After introducing notation in Section 4.1, we analyze several particular choices of $\mathcal{L}_2$ and $\mathbf{b}$ to build our understanding of the set $\mathcal{S}$. In Section 4.2, we look at a bijective linear operator $\mathcal{L}_2$ mapping the sum of the squared polynomials into the vector of its coefficients. This will lead us to our first¹ *embedded submanifold* of $\mathbb{R}[x]_d^2$. Next, in Section 4.3, we show that the restriction of $\mathcal{S}$ to pairs of coprime polynomials is also an embedded submanifold of $\mathbb{R}[x]_d^2$, independently of the choice of $\mathcal{L}_2$ and $\mathbf{b}$ (assuming $\mathcal{S}$ is not empty).

However, for generic $\mathcal{L}_2$ and $\mathbf{b}$, $\mathcal{S}$ is not always an embedded submanifold of $\mathbb{R}[x]_d^2$. This will be explained in more detail in the [next chapter](#).

4.1 Notations

To simplify the notation, given two polynomials $u_1, u_2 \in \mathbb{R}[x]_d$, we define the vector $\mathbf{u} = (u_1, u_2) \in \mathbb{R}[x]_d^2$ and the sum-of-squares operator

$$\sigma : \mathbb{R}[x]_d^2 \rightarrow \mathbb{R}[x]_{2d} : \mathbf{u} \mapsto u_1^2 + u_2^2.$$

The differential of σ at a point $\mathbf{u} \in \mathbb{R}[x]_d^2$ in a direction $\mathbf{v} \in \mathbb{R}[x]_d^2$ is given by

$$D\sigma(\mathbf{u})[\mathbf{v}] := \lim_{t \rightarrow 0} \frac{\sigma(\mathbf{u} + t\mathbf{v}) - \sigma(\mathbf{u})}{t} = 2(u_1v_1 + u_2v_2).$$

¹The manifold of pairs of coprime degree- d polynomials presented in Chapter 3 is also an embedded submanifold of $\mathbb{R}[x]_d^2$, but by first embedded submanifold we mean the first example with dimension strictly less than the embedding space.

Correspondingly, we define the (linear) operator

$$\mathcal{A}_u : \mathbb{R}[x]_d^2 \rightarrow \mathbb{R}[x]_{2d} : \mathbf{v} \mapsto u_1 v_1 + u_2 v_2.$$

Lemma 4.1

Let $\mathbf{u} \in \mathbb{R}[x]_d^2$, and let $g = \gcd(u_1, u_2)$ be the greatest common divisor between the two polynomials of $\mathbf{u}$. If g has degree K , then $\text{Im } \mathcal{A}_u = g \cdot \mathbb{R}[x]_{2d-K}$ is of dimension $2d - K + 1$, and the kernel of $\mathcal{A}_u$ has dimension $K + 1$.

Proof. By definition of $\mathcal{A}_u$, if g divides both u_1 and u_2 , then it will divide $\mathcal{A}_u(\mathbf{v})$ for every $\mathbf{v} \in \mathbb{R}[x]_d^2$. Therefore, the operator $\mathcal{A}_u$ is equal to $g \cdot \mathcal{A}_{\bar{\mathbf{u}}}$, where $\bar{\mathbf{u}}$ is defined such that $\mathbf{u} = g \cdot \bar{\mathbf{u}}$ (i.e., it contains the coprime factors of u_1 and u_2). Note that $\mathcal{A}_{\bar{\mathbf{u}}}$ still takes as input a pair of degree- d polynomials, even though the polynomials $\bar{u}_1$ and $\bar{u}_2$ are of degree $d - K$.

The rest of the proof follows from [LYP23, Lemma 3.3], in which the authors prove that the image of $\mathcal{A}_{\bar{\mathbf{u}}}$ is $\mathbb{R}[x]_{2d-K}$, $\bar{u}_1$ and $\bar{u}_2$ being coprime. Therefore, the image of $\mathcal{A}_u$ is of dimension $2d - K + 1$, and its kernel is of dimension $K + 1$ by the rank-nullity theorem. $\square$

Remark 4.2. The matrix associated to the linear operator $\mathcal{A}_u$, when expressed in the monomial basis, is similar to the Sylvester matrix (Definition 3.3). $\triangle$

4.2 Embedded submanifold with bijective operator

Consider a problem in which we seek to optimize some function on the set of sum-of-squares decompositions of a polynomial $p \in \mathbb{R}[x]_{2d}$. This corresponds to the problem (1.2) whose feasible set $\mathcal{S}$ (4.1) is defined by an operator $\mathcal{L}_2$ mapping a polynomial in $\mathbb{R}[x]_{2d}$ into its vector of coefficients in $\mathbb{R}^{2d+1}$ (considering the monomial basis), and a constraint vector $\mathbf{b}$ associated to a nonnegative polynomial (the set $\mathcal{S}$ would be empty otherwise). As shown in the proof of Lemma 2.2, $\mathcal{S}$ is the union of at most 2^{K-1} “circles” corresponding to different SOS decompositions that may be rotated. We show in Theorem 4.4 that $\mathcal{S}$ is an *embedded submanifold* [Bou23, Theorem 8.75] of $\mathbb{R}[x]_d^2$ of dimension 1.

Definition 4.3 – Embedded submanifold

Let $\mathcal{M}$ be a smooth manifold of dimension n . A non-empty subset $\mathcal{N}$ of $\mathcal{M}$ is a smooth *embedded submanifold* of $\mathcal{M}$ of dimension $n - k$ if and only if, for a fixed integer $k \geq 1$ and for each $x \in \mathcal{N}$, there exists a neighbourhood $\mathcal{U}$ of x in $\mathcal{M}$ and a smooth function $h : \mathcal{U} \rightarrow \mathbb{R}^k$ such that:

1. If y is in $\mathcal{U}$, then $h(y) = 0$ if and only if $y \in \mathcal{N}$;
2. The rank of the differential of h at x , $Dh(x)$, is equal to k .

Such a function h is called a *local defining function* for $\mathcal{N}$ at x , and $\mathcal{M}$ is called the *embedding space*.

In Definition 4.3, we use the concept of *differential* of a smooth map between two manifolds. While the exact definition ([Bou23, Def 8.37]) involves the *tangent spaces* of the

manifolds at the considered points and a *smooth curve* on $\mathcal{M}$ satisfying some conditions, we can keep it simple, since the input space of the local defining function h we will consider is an open submanifold of a vector space, and its output space is a linear space. In that case, the differential of h is defined in the classical way as

$$Dh(\mathbf{u}) : \mathcal{E} \rightarrow \mathbb{R}^k : \mathbf{v} \mapsto \lim_{t \rightarrow 0} \frac{h(\mathbf{u} + t\mathbf{v}) - h(\mathbf{u})}{t},$$

where $\mathcal{E}$ is the vector space in which the open set $\mathcal{U}$ is contained. The differential is a linear operator, and its *rank* is equal to the dimension of its image.

Theorem 4.4

Let $p \in \mathbb{R}[x]_{2d}$ be a nonnegative polynomial. The set $\mathcal{S} := \{\mathbf{u} \in \mathbb{R}[x]_d^2 \mid \sigma(\mathbf{u}) = p\}$ of sum-of-squares decompositions of p is an embedded submanifold of $\mathbb{R}[x]_d^2$, of dimension 1, corresponding to the union of at most 2^{K-1} disjoint connected components.

Sketch of the proof. To show this result, we proceed in three steps. In Lemma 4.5, we show that the different circles of the SOS decompositions obtained in the proof of Lemma 2.2 do not intersect. The definition of embedded submanifold mentions that, for each point of the set, there exists a neighbourhood and a smooth local defining function. These may differ from one point to another. In particular, as the circles do not intersect, there exist neighbourhoods around each circle that do not intersect². Therefore, we may consider one circle at a time. If these are all embedded submanifolds of $\mathbb{R}[x]_d^2$ of the same dimension, their union will be.

Next, in Lemma 4.6, we show that the common roots between the two polynomials of a decomposition stay the same along the circle. This allows us to consider the same circle in a lower-dimensional space, in which the polynomials along the circle are coprime, by factorizing the common factors. We show in Lemma 4.7 that this circle is a one-dimensional embedded submanifold of $\mathbb{R}[x]_{d-K}^2$, where K is the number of shared roots between the polynomials along the circle. We can therefore conclude by saying that the original circle is itself a one-dimensional embedded submanifold of $\mathbb{R}[x]_d^2$, since the mapping consisting of multiplying the lower-dimensional circle by the polynomial of the shared roots is a smooth diffeomorphic embedding (i.e., a bijective map that is smooth and with a smooth inverse — see [Lee12, Proposition 5.2])³. $\square$

As described in the sketch of the proof of Theorem 4.4, we start by showing that the circles corresponding to two distinct SOS decompositions of a nonnegative polynomial $p \in \mathbb{R}[x]_{2d}$ do not intersect.

Lemma 4.5

Let $(u_1, u_2) \in \mathbb{R}[x]_d^2$ and $(v_1, v_2) \in \mathbb{R}[x]_d^2$ be two SOS decompositions of p . Let $R(\cdot)$ denote a 2×2 rotation matrix. If $R(\theta) \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = R(\phi) \begin{pmatrix} v_1 \\ v_2 \end{pmatrix}$ for some $\theta, \phi \in [0, 2\pi[$, the two points belong to the same circle, i.e., $\begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = R(\psi) \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}$ for some $\psi \in [0, 2\pi[$.

²This follows from the fact that the topology defined in Section 3.1 is Hausdorff.

³Intuitively, multiplying by a fixed polynomial should preserve the smooth manifold structure.

Proof. This follows from the regular character of rotation matrices. We have

$$\begin{pmatrix} v_1 \\ v_2 \end{pmatrix} = R(\phi)^{-1}R(\theta) \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = R(\theta - \phi) \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}.$$

□

Lemma 4.6

Let $(u_1, u_2) \in \mathbb{R}[x]_d^2$ be an SOS decomposition of p , and let $g := \gcd(u_1, u_2)$ be their greatest common divisor. Defining the polynomials $\bar{u}_1, \bar{u}_2 \in \mathbb{R}[x]_{d-K}$ such that $u_1 = g \cdot \bar{u}_1$ and $u_2 = g \cdot \bar{u}_2$, where K is the number of shared roots, the resultant of $\bar{u}_1$ and $\bar{u}_2$ is constant along the circle $R(\theta) \begin{pmatrix} \bar{u}_1 \\ \bar{u}_2 \end{pmatrix}$, $\theta \in [0, 2\pi[$.

Proof. Let $\bar{u}_1 = \sum_{i=0}^m a_i x^i$ and $\bar{u}_2 = \sum_{i=0}^n b_i x^i$, with m and n the degrees of $\bar{u}_1$ and $\bar{u}_2$ respectively. For simplicity, assume $m = n$, but a similar argument can be made for $m \neq n$.

The Sylvester matrix of $R(\theta) \begin{pmatrix} \bar{u}_1 \\ \bar{u}_2 \end{pmatrix}$ is

$$S := \text{Syl} \left(\begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} \bar{u}_1 \\ \bar{u}_2 \end{pmatrix} \right) = \begin{pmatrix} \cos(\theta)a_m - \sin(\theta)b_m & & & \sin(\theta)a_m + \cos(\theta)b_m & & \\ \vdots & \ddots & & \vdots & \ddots & \\ \cos(\theta)a_1 - \sin(\theta)b_1 & \dots & \cos(\theta)a_m - \sin(\theta)b_m & \sin(\theta)a_1 + \cos(\theta)b_1 & \dots & \sin(\theta)a_m + \cos(\theta)b_m \\ \cos(\theta)a_0 - \sin(\theta)b_0 & \dots & \cos(\theta)a_{m-1} - \sin(\theta)b_{m-1} & \sin(\theta)a_0 + \cos(\theta)b_0 & \dots & \sin(\theta)a_{m-1} + \cos(\theta)b_{m-1} \\ & \ddots & \vdots & & \ddots & \vdots \\ & & \cos(\theta)a_0 - \sin(\theta)b_0 & & & \sin(\theta)a_0 + \cos(\theta)b_0 \end{pmatrix}.$$

Since the determinant is multilinear in its columns [Str16, Section 5.1], the determinant of the Sylvester matrix is

$$\det(S) = \cos^2(\theta) \begin{vmatrix} a_m & & b_m & \\ \vdots & & \vdots & \\ a_0 & S[2:m] & b_0 & S[m+2:2m] \end{vmatrix} - \sin^2(\theta) \begin{vmatrix} b_m & & a_m & \\ \vdots & & \vdots & \\ b_0 & S[2:m] & a_0 & S[m+2:2m] \end{vmatrix}.$$

By swapping the two isolated columns of the second matrix, the determinant is multiplied by -1 . Using the Pythagorean identity, the determinant of the Sylvester matrix is

$$\det(S) = \begin{vmatrix} a_m & & b_m & \\ \vdots & & \vdots & \\ a_0 & S[2:m] & b_0 & S[m+2:2m] \end{vmatrix}.$$

Applying the same reasoning to the $2(m-1)$ other columns, we get that

$$\det \left(\text{Syl} \left(\begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} \bar{u}_1 \\ \bar{u}_2 \end{pmatrix} \right) \right) = \det(\text{Syl}(\bar{u}_1, \bar{u}_2)).$$

□

Lemma 4.6 implies that the common roots between the two polynomials are identical over the whole circle. Indeed, if $u_1 = g\bar{u}_1$ and $u_2 = g\bar{u}_2$, we have

$$R(\theta) \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} = gR(\theta) \begin{pmatrix} \bar{u}_1 \\ \bar{u}_2 \end{pmatrix}.$$

Since the resultant is constant over the whole circle, there is no point on the circle where the two polynomials u_1 and u_2 may share a root different than the ones contained in g .

Lemma 4.7

If $\mathbf{u} = (u_1, u_2) \in \mathbb{R}[x]_d^{2,c}$ contains two coprime polynomials of degree d , then the circle defined by $\mathcal{C} = R(\theta) \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}$, $\theta \in [0, 2\pi[$ defines an embedded submanifold of dimension 1 of $\mathbb{R}[x]_d^2$.

Proof. According to Definition 4.3, we consider as candidate local defining function $h : \mathcal{U} \rightarrow \mathbb{R}[x]_{2d} : \mathbf{u} \mapsto \sigma(\mathbf{u}) - p$, where $\mathcal{U}$ is a neighbourhood of $\mathcal{C}$ in $\mathbb{R}[x]_d^2$. If $\mathcal{U}$ is chosen to not contain any other circle, a point $\mathbf{u}$ belongs to $\mathcal{C}$ if and only if it is a zero of h . Furthermore, the differential of h at a point $\mathbf{u} \in \mathcal{C}$ is given by

$$Dh(\mathbf{u}) : \mathbb{R}[x]_d^2 \rightarrow \mathbb{R}[x]_{2d} : \mathbf{v} \mapsto \mathcal{A}_{\mathbf{u}}(\mathbf{v}),$$

which, by Lemma 4.1, is of rank $2d + 1$ independently of $\mathbf{u} \in \mathcal{C}$, since u_1 and u_2 are coprime. Therefore, $\mathcal{C}$ is an embedded submanifold of $\mathbb{R}[x]_d^2$ of dimension

$$\dim \mathbb{R}[x]_d^2 - \dim \mathbb{R}[x]_{2d} = 1.$$

□

This result concludes the proof of Theorem 4.4.

Remark 4.8. A similar reasoning shows that $\mathcal{S}$ is an embedded submanifold of $\mathbb{R}[x]_d^2$ for every bijective operator $\mathcal{L}_2$. $\triangle$

We will conclude this section by briefly discussing what Theorem 4.4 implies on the main problem (1.2) we want to solve. We have shown that, if $\mathcal{L}_2$ is a bijective operator, the feasible set $\mathcal{S}$ is a smooth manifold. A first-order optimization algorithm for embedded submanifolds of $\mathbb{R}[x]_d^2$ is presented in Chapter 6. However, as we have shown, when the constraint vector $\mathbf{b}$ corresponds to a polynomial with at least two distinct pairs of complex conjugate roots, the manifold $\mathcal{S}$ is a union of disjoint circles. This is problematic: how could the algorithm jump from one circle to another? Furthermore, if we apply the algorithm of Legat et al. [LYP23] presented in Section 2.4 to find a first feasible point, we cannot predict on which circle it will be. In fact, this is not an issue at all! Indeed, the objective of (1.2) concerns $\sigma(\mathbf{u})$, which is fully determined when the constraint operator $\mathcal{L}_2$ is bijective: $\sigma(\mathbf{u}) = \mathcal{L}_2^{-1}(\mathbf{b})$. Therefore, each feasible point $\mathbf{u} \in \mathcal{S}$ has the same objective value, and it is not necessary to apply an optimization algorithm.

4.3 Embedded submanifolds of coprime polynomials

As mentioned at the end of the previous section, the feasible set $\mathcal{S}$ is a smooth manifold when $\mathcal{L}_2$ is bijective, but this is useless as we do not need to optimize on the feasible set: the objective is constant on it. In this section, we consider a restriction $\bar{\mathcal{S}}$ of $\mathcal{S}$ to the pairs of coprime polynomials $\mathbb{R}[x]_d^{2,c}$, and show that $\bar{\mathcal{S}}$ is an embedded submanifold of $\mathbb{R}[x]_d^{2,c}$. $\mathbb{R}[x]_d^{2,c}$ being itself embedded in $\mathbb{R}[x]_d^2$, $\bar{\mathcal{S}}$ is an embedded submanifold of $\mathbb{R}[x]_d^2$.

Theorem 4.9

The restriction of $\mathcal{S}$ to the set of coprime polynomials,

$$\bar{\mathcal{S}} := \left\{ (u_1, u_2) \in \mathbb{R}[x]_d^{2,c} \mid \mathcal{L}_2 \left(u_1^2(x) + u_2^2(x) \right) = \mathbf{b} \right\},$$

where $\mathcal{L}_2 : \mathbb{R}[x]_{2d} \rightarrow \mathbb{R}^m$ is surjective and $\mathbf{b} \in \mathbb{R}^m$, is an embedded submanifold of $\mathbb{R}[x]_d^{2,c}$ of dimension $2d + 2 - m$.

Proof. Consider the candidate local defining function $h : \mathbb{R}[x]_d^{2,c} \rightarrow \mathbb{R}^m : \mathbf{u} \mapsto \mathcal{L}_2(\sigma(\mathbf{u})) - \mathbf{b}$. We can easily check that h vanishes exactly at the points of $\bar{\mathcal{S}}$. Furthermore, the differential of h at a point $\mathbf{u} \in \bar{\mathcal{S}}$ is

$$Dh(\mathbf{u}) : \mathbb{R}[x]_d^2 \rightarrow \mathbb{R}^m : \mathbf{v} \mapsto 2\mathcal{L}_2 \circ \mathcal{A}_{\mathbf{u}}(\mathbf{v}).$$

The polynomials u_1 and u_2 being coprime, $\mathcal{A}_{\mathbf{u}}$ is surjective (Lemma 4.1), and since $\mathcal{L}_2$ is also surjective, so is their composition. The application $Dh(\mathbf{u})$ is therefore of full rank at every point of $\bar{\mathcal{S}}$. As the dimensions of the input and output spaces of $Dh(\mathbf{u})$ are $2d + 2$ and m respectively, we can conclude on the dimension of the smooth embedded submanifold. $\square$

Remark 4.10. If all the points in the set $\mathcal{S}$ are coprime, Theorem 4.9 tells us that $\mathcal{S}$ is a smooth embedded submanifold of $\mathbb{R}[x]_d^2$. In the same way, if all the elements of $\mathcal{S}$ are coprime pairs in $\mathbb{R}[x]_{d-l}^2$ multiplied by the same polynomial $g \in \mathbb{R}[x]_l$, $\mathcal{S}$ is a smooth manifold as it is a smooth embedded submanifold of a lower-dimensional space ($\mathbb{R}[x]_{d-l}^2$), embedded in $\mathbb{R}[x]_d^2$ through a multiplication by the fixed polynomial g (which is a smooth diffeomorphic embedding like in the proof of Theorem 4.4). Yet, there is no general technique for determining whether $\mathcal{S}$ has such a structure. $\triangle$

Chapter 5

Stratification structure on SOS decompositions of polynomials

Inspired by the [last chapter](#), a natural question to ask now is whether the set $\mathcal{S}$ of (4.1),

$$\mathcal{S} := \left\{ (u_1, u_2) \in \mathbb{R}[x]_d^2 \mid \mathcal{L}_2 \left(u_1^2(x) + u_2^2(x) \right) = \mathbf{b} \right\},$$

is an embedded submanifold of $\mathbb{R}[x]_d^2$ for every linear operator $\mathcal{L}_2$ and constraint vector $\mathbf{b}$. Unfortunately, the answer is no. The aim of this chapter is to explain the reasons for this failure, and to suggest possible solutions.

In Section 5.1, we design a counterexample for which the function $h : \mathbb{R}[x]_d^2 \rightarrow \mathbb{R}^m : \mathbf{u} \mapsto \mathcal{L}_2(\sigma(\mathbf{u})) - \mathbf{b}$ is not a local defining function. Nevertheless, this simply means that the particular candidate local defining function h considered is not valid. In Section 5.2, we derive an expression of the *tangent cone* of $\mathcal{S}$ at each of its points. We then examine the cone in detail to determine in which case the set $\mathcal{S}$ is not a manifold. However, in Section 5.3, we show for a particular example that it is a *stratified set*. Though stratified sets are not as nice as smooth manifolds, algorithms have recently been developed tackling the issues they present. It will be interesting in future work to analyze those algorithms, now that we know we face stratified sets.

5.1 A failure for our candidate local defining function

In the previous chapter, we frequently used the function

$$h : \mathcal{U} \rightarrow \mathbb{R}^m : \mathbf{u} \mapsto \mathcal{L}_2(\sigma(\mathbf{u})) - \mathbf{b}, \tag{5.1}$$

for some open subsets $\mathcal{U}$ of $\mathbb{R}[x]_d^2$, as a local defining function. The question is now whether we can use it to show that $\mathcal{S}$ is an embedded submanifold of $\mathbb{R}[x]_d^2$. It is easily shown that, for every open subset $\mathcal{U}$ of $\mathbb{R}[x]_d^2$, $h^{-1}(\mathbf{0}) = \mathcal{U} \cap \mathcal{S}$, where

$$h^{-1}(\mathbf{0}) = \left\{ \mathbf{u} \in \mathbb{R}[x]_d^2 \mid h(\mathbf{u}) = \mathbf{0} \right\}.$$

Therefore, we may focus on the question of the rank of the differential of h .

At a point $\mathbf{u} \in \mathcal{S}$, the differential of h is equal to (following Section 4.1)

$$Dh(\mathbf{u}) : \mathbb{R}[x]_d^2 \rightarrow \mathbb{R}^m : \mathbf{v} \mapsto 2\mathcal{L}_2 \circ \mathcal{A}_{\mathbf{u}}(\mathbf{v}).$$

Lemma 5.1

The kernel of $Dh(\mathbf{u}) = 2\mathcal{L}_2 \circ \mathcal{A}_u$ is given by

$$\ker Dh(\mathbf{u}) = \mathcal{A}_u^{-1}(\ker \mathcal{L}_2) \quad (5.2a)$$

$$= \ker \mathcal{A}_u \oplus \mathcal{A}_u^\dagger(\text{Im } \mathcal{A}_u \cap \ker \mathcal{L}_2), \quad (5.2b)$$

where $\mathcal{A}_u^{-1}(\mathcal{X})$ denotes the preimage of a set $\mathcal{X} \subseteq \mathbb{R}[x]_{2d}$ through the operator $\mathcal{A}_u$, and $\mathcal{A}_u^\dagger$ denotes the *Moore-Penrose pseudo-inverse* of the operator $\mathcal{A}_u$ (see [BH12] for more information).

Proof. For the first equality, by definition of the preimage of a linear application,

$$\begin{aligned} \ker Dh(\mathbf{u}) &= \left\{ \mathbf{v} \in \mathbb{R}[x]_d^2 \mid (\mathcal{L}_2 \circ \mathcal{A}_u)(\mathbf{v}) = \mathbf{0} \right\} \\ &= \left\{ \mathbf{v} \in \mathbb{R}[x]_d^2 \mid \mathcal{A}_u(\mathbf{v}) \in \ker \mathcal{L}_2 \right\} \\ &= \mathcal{A}_u^{-1}(\ker \mathcal{L}_2). \end{aligned}$$

For the second equality, we prove the mutual inclusion.

$\subseteq$ Let $\mathbf{v} \in \mathcal{A}_u^{-1}(\ker \mathcal{L}_2)$. Then, $\mathcal{A}_u(\mathbf{v}) \in \text{Im } \mathcal{A}_u \cap \ker \mathcal{L}_2$. By [BH12, Theorem 6.1], this implies that $\mathbf{v} \in \ker \mathcal{A}_u \oplus \mathcal{A}_u^\dagger(\text{Im } \mathcal{A}_u \cap \ker \mathcal{L}_2)$.

$\supseteq$ Let $\mathbf{v} \in \ker \mathcal{A}_u \oplus \mathcal{A}_u^\dagger(\text{Im } \mathcal{A}_u \cap \ker \mathcal{L}_2)$. Then, there exists $\mathbf{w}_1 \in \ker \mathcal{A}_u$ and $\mathbf{w}_2 \in \mathcal{A}_u^\dagger(\text{Im } \mathcal{A}_u \cap \ker \mathcal{L}_2)$ such that $\mathbf{v} = \alpha \mathbf{w}_1 + \beta \mathbf{w}_2$ for some coefficients $\alpha, \beta \in \mathbb{R}$. This implies that

$$(\mathcal{L}_2 \circ \mathcal{A}_u)(\mathbf{v}) = \beta(\mathcal{L}_2 \circ \mathcal{A}_u)(\mathbf{w}_2) = \beta \mathcal{L}_2(\mathcal{A}_u(\mathbf{w}_2)) = \mathbf{0}.$$

The last equality follows from the fact that, by definition of $\mathbf{w}_2$, there exists $\mathbf{w}_3 \in \mathbb{R}[x]_d^2$ such that $\mathcal{A}_u(\mathbf{w}_3) \in \ker \mathcal{L}_2$ and $\mathbf{w}_2 = \mathcal{A}_u^\dagger(\mathcal{A}_u(\mathbf{w}_3))$ (the equality is due to the uniqueness of the Moore-Penrose pseudo-inverse [BH12, Section 2.3]). By definition of the Moore-Penrose pseudo-inverse ([BH12, Section 2.2]),

$$\mathcal{L}_2(\mathcal{A}_u(\mathbf{w}_2)) = \mathcal{L}_2 \left(\mathcal{A}_u \left(\mathcal{A}_u^\dagger(\mathcal{A}_u(\mathbf{w}_3)) \right) \right) = \mathcal{L}_2(\mathcal{A}_u(\mathbf{w}_3)) = \mathbf{0}.$$

□

5.1.1 Communicating vessels

Using the second expression (5.2b) for the kernel of $Dh(\mathbf{u})$, for h to be a local defining function, any increase or decrease in the dimension of the kernel of $\mathcal{A}_u$ as $\mathbf{u}$ varies must be compensated by a corresponding decrease or increase in the dimension of $\mathcal{A}_u^\dagger(\text{Im } \mathcal{A}_u \cap \ker \mathcal{L}_2)$. This can be seen as communicating vessels, as the sum of the two dimensions must remain constant.

The dimension of the kernel of $\mathcal{A}_u$ only depends on the number K of common roots between the two polynomials of $\mathbf{u}$, and is given by $K + 1$ (Lemma 4.1). Regarding the

second term $\mathcal{A}_u^\dagger(\text{Im } \mathcal{A}_u \cap \ker \mathcal{L}_2)$, its dimension is equal ¹ to the dimension of $\text{Im } \mathcal{A}_u \cap \ker \mathcal{L}_2$. When the dimension of the kernel of $\mathcal{A}_u$ changes, that of the image of $\mathcal{A}_u$ changes by the same amount in the opposite direction, by the rank-nullity theorem. However, problems arise if the dimension decrease of $\text{Im } \mathcal{A}_u$ does not take place within the kernel of $\mathcal{L}_2$, in which case the dimension of their intersection remains unchanged.

In the following example, we analyze the case of a set $\mathcal{S}$ containing points which are coprime polynomials, as well as pairs of polynomials with exactly one common root.

Example 5.2. Consider a pair of coprime polynomials $\mathbf{u}$ that belongs to $\mathcal{S}$. At that point, the operator $\mathcal{A}_u$ is surjective (Lemma 4.1). Its kernel is therefore of dimension one; and the intersection of the kernel of $\mathcal{L}_2$ with the image of $\mathcal{A}_u$ is the kernel of $\mathcal{L}_2$, of dimension $2d + 1 - m$. Hence, the kernel of $Dh(\mathbf{u})$ has dimension $2d + 2 - m$.

Now, consider a pair of polynomials $\mathbf{v}$ sharing one common root. The dimension of the kernel of $\mathcal{A}_v$ is one unit greater. For the rank of the differential to stay constant, the intersection of $\text{Im } \mathcal{A}_v$ and $\ker \mathcal{L}_2$ has to decrease by one correspondingly. Yet, if every element in the kernel of $\mathcal{L}_2$ vanishes at the common root, then $\ker \mathcal{L}_2$ is contained in the image of $\mathcal{A}_v$, and their intersection keeps the same dimension (as it is still $\ker \mathcal{L}_2$). In that case, the rank of the differential increases by one, which shows that h is not a valid local defining function. $\diamond$

Remark 5.3. There is a way around the issue observed in Example 5.2. If $\ker \mathcal{L}_2 \subseteq \text{Im } \mathcal{A}_v$, there can not exist a pair of coprime polynomials in the set $\mathcal{S}$. By contradiction, assume $\tilde{\mathbf{u}}$ is a pair of coprime polynomials belonging to the set $\mathcal{S}$. Then, the difference between its sum of squares and the sum of squares of $\mathbf{v}$ must belong to the kernel of $\mathcal{L}_2$:

$$\mathcal{L}_2(\sigma(\mathbf{v})) = \mathbf{b} = \mathcal{L}_2(\sigma(\tilde{\mathbf{u}})) \Rightarrow \sigma(\mathbf{v}) - \sigma(\tilde{\mathbf{u}}) \in \ker \mathcal{L}_2.$$

However, since the kernel of $\mathcal{L}_2$ is included in the image of $\mathcal{A}_v$, every element in the kernel of $\mathcal{L}_2$ must vanish at the common roots of the polynomials in $\mathbf{v}$. As $\sigma(\mathbf{v})$ vanishes at those roots, $\sigma(\tilde{\mathbf{u}})$ must also vanish, and therefore shares the same roots as $\mathbf{v}$. $\triangle$

Remark 5.3 means that, as long as a pair of polynomials $\mathbf{u}$ belonging to the set $\mathcal{S}$ is such that $\ker \mathcal{L}_2 \subseteq \text{Im } \mathcal{A}_u$, all the elements in the set $\mathcal{S}$ have to share the same common roots² as the ones of the two polynomials of $\mathbf{u}$. For Example 5.2, this means that if the dimension of the kernel of $Dh(\cdot)$ changes at a noncoprime pair $\mathbf{v}$ (with one common root), there is in fact no coprime pair in $\mathcal{S}$. Therefore, the dimension does not really change, it is just different to the case of having coprime polynomials in $\mathcal{S}$.

5.1.2 A counterexample

Therefore, we have to look further, for an example in which the kernel of $\mathcal{L}_2$ is not contained in the image of $\mathcal{A}_v$. We will construct this example step by step. For simplicity,

¹This follows from the injectivity of $\mathcal{A}_u^\dagger$ when restricted to $\text{Im } \mathcal{A}_u$. Indeed, by contradiction, assume $p \in \text{Im } \mathcal{A}_u$ is nonzero and belongs to the kernel of $\mathcal{A}_u^\dagger$. By [BH12, Theorem 6.1], $\mathcal{A}_u^{-1}(p) = \mathcal{A}_u^\dagger(p) + \ker \mathcal{A}_u = \ker \mathcal{A}_u$. This would imply that $p = 0$, in contradiction with its definition.

²This holds including the multiplicity of the roots. Indeed, if a root is of multiplicity $\mu > 1$, then the $\mu - 1$ first derivatives of all polynomials in $\ker \mathcal{L}_2$ vanish at that root. The same reasoning as the one used in Remark 5.3 allows to conclude that the polynomials of $\tilde{\mathbf{u}}$ must share all the common roots as the one of $\mathbf{v}$, including their multiplicity.

we consider an application $\mathcal{L}_2 : \mathbb{R}[x]_{2d} \rightarrow \mathbb{R}$ (with a scalar output), with a $2d$ -dimensional kernel (included in the $2d + 1$ -dimensional space $\mathbb{R}[x]_{2d+1}$).

Assume $\mathbf{u}$ is a pair of coprime polynomials that belongs to $\mathcal{S}$. Then, the kernel of $Dh(\mathbf{u})$ has dimension $2d + 1$. Now, if the two polynomials share one common root, either $\ker \mathcal{L}_2 = \text{Im } \mathcal{A}_{\mathbf{u}}$, in which case Remark 5.3 tells us that no coprime polynomials can belong to $\mathcal{S}$, or the dimension of the intersection between the kernel of $\mathcal{L}_2$ and the image of $\mathcal{A}_{\mathbf{u}}$ decreases by one, compensating the increase of the dimension of the kernel of $\mathcal{A}_{\mathbf{u}}$.

Thus, we have to look at a pair of polynomials sharing two roots. Let $\mathbf{v}$ be such a pair in $\mathcal{S}$. Then, the dimension³ of the intersection between $\ker \mathcal{L}_2$ and $\text{Im } \mathcal{A}_{\mathbf{v}}$ is either $2d - 1$ (which is the dimension of $\text{Im } \mathcal{A}_{\mathbf{v}}$), or $2d - 2$. The second case is of no interest to us, as it consists of keeping the same rank for $Dh(\mathbf{v})$ (as the dimension of $\ker \mathcal{A}_{\mathbf{v}}$ is 3). Yet, if the dimension of the intersection is $2d - 1$, the rank of $Dh(\mathbf{v})$ is 3 (from the kernel of $\mathcal{A}_{\mathbf{v}}$) plus $2d - 1$, giving $2d + 2$. It has therefore increased by one.

Remark 5.4. If the intersection is of dimension $2d - 1$, then the intersection is equal to $\text{Im } \mathcal{A}_{\mathbf{v}}$, or in other words $\text{Im } \mathcal{A}_{\mathbf{v}} \subset \ker \mathcal{L}_2$. Since $\mathbf{v} \in \mathcal{S}$, the constraint vector (which is a scalar in this case) $\mathbf{b}$ is equal to $\mathcal{L}_2(\sigma(\mathbf{v})) = 0$. In order to have a pair of coprime polynomials $\mathbf{u}$ in the set $\mathcal{S}$, the last direction⁴ of the kernel of $\mathcal{L}_2$ has to be $\sigma(\mathbf{u})$. $\triangle$

Example 5.5. It is interesting to consider a numerical example at this point. We will consider pairs of polynomials of degree 2 as we require at least two common roots. Assume the pair of polynomials $\mathbf{v} = ((x - 1)(x - 3), 2(x - 1)(x - 3))$ belongs to the set $\mathcal{S}$. A possible choice for this is to consider a linear operator

$$\mathcal{L}_2 : \mathbb{R}[x]_4 \rightarrow \mathbb{R} : p \mapsto p(3) - p(1).$$

With this choice, the image of $\mathcal{A}_{\mathbf{v}}$ is included in the kernel of $\mathcal{L}_2$, as every element in the image of $\mathcal{A}_{\mathbf{v}}$ vanishes at 1 and at 3. This is already a 3-dimensional subset of the kernel of $\mathcal{L}_2$, which is 4-dimensional. A last direction is given by any polynomial p such that $p(1) = p(3) \neq 0$. For example, the pair of polynomials $\mathbf{u} = (x(x - 2), -x^2 + 3x + 1)$ belongs to the set $\mathcal{S}$ as its sum of squares is evaluated at the same (nonzero) value at 1 and at 3. Its sum of squares gives therefore a last direction for the kernel of $\mathcal{L}_2$. At $\mathbf{u}$, the dimension of the kernel of $Dh(\cdot)$ is 5 (since $\mathbf{u}$ is a pair of coprime polynomials), while it is 6 at $\mathbf{v}$ (since $\text{Im } \mathcal{A}_{\mathbf{v}}$ is a subset of $\ker \mathcal{L}_2$ and the polynomials of $\mathbf{v}$ share two roots). $\diamond$

5.2 Tangent cone

At this stage, Example 5.5 tells us only that h is not a local defining function in a neighbourhood of $\mathbf{v}$ in $\mathbb{R}[x]_d^2$. Yet, the set $\mathcal{S}$ described in Example 5.5 could still be a smooth manifold. To show that this is not the case, we rely on a necessary condition which states that the *tangent cone* at each point of the manifold must be a vector space.

In Section 5.2.1, we first derive an implicit expression for the *tangent cone* [RW98; Oli+25] to the set $\mathcal{S}$ at each of its points. In Section 5.2.2, we will then analyze the

³It can not be more as it is bounded by the minimal dimension of the intersecting sets. Also, it can not be less since $\text{Im } \mathcal{A}_{\mathbf{v}} + \ker \mathcal{L}_2 = \text{Im } \mathcal{A}_{\mathbf{v}} \setminus (\text{Im } \mathcal{A}_{\mathbf{v}} \cap \ker \mathcal{L}_2) \oplus (\text{Im } \mathcal{A}_{\mathbf{v}} \cap \ker \mathcal{L}_2) \oplus \ker \mathcal{L}_2 \setminus (\text{Im } \mathcal{A}_{\mathbf{v}} \cap \ker \mathcal{L}_2) \subseteq \mathbb{R}[x]_{2d}$, which is of dimension $2d + 1$.

⁴As the kernel of $\mathcal{L}_2$ is $2d$ -dimensional and there are already $2d - 1$ directions filled in by the image of $\mathcal{A}_{\mathbf{v}}$.

equations (5.4) describing the tangent cone. In Section 5.2.3, we show that the tangent cone on the smooth embedded submanifolds of Chapter 4 is indeed linear. Yet, for the set $\mathcal{S}$ described in Example 5.5, the tangent cone is not linear, as it is explained in Section 5.2.4.

5.2.1 Expression of the tangent cone

Definition 5.6 – Tangent vector

A vector $\mathbf{v} \in \mathbb{R}[x]_d^2$ is said to be *tangent* to $\mathcal{S}$ at $\mathbf{u} \in \mathcal{S}$ if there exist sequences $(\mathbf{v}_i)_{i \in \mathbb{N}}$ in $\mathcal{S}$ converging to $\mathbf{u}$ and $(t_i)_{i \in \mathbb{N}}$ in $(0, \infty)$ such that the sequence $\left(\frac{\mathbf{v}_i - \mathbf{u}}{t_i}\right)_{i \in \mathbb{N}}$ converges to $\mathbf{v}$.

Definition 5.7 – Tangent cone

The *tangent cone*^a of $\mathcal{S}$ at a point $\mathbf{u} \in \mathcal{S}$ is the set of tangent vectors of $\mathcal{S}$ at $\mathbf{u}$.

^aThis is a cone in the usual sense, i.e., if $\mathbf{v}$ belongs to the tangent cone, then $\lambda \mathbf{v}$ belongs to the tangent cone for all $\lambda \geq 0$. However, it may not be a convex cone.

To get an expression of the tangent cone at a point $\mathbf{u} \in \mathcal{S}$, we will find the set of vectors $\mathbf{v} \in \mathbb{R}[x]_d^2$ for which there exists such sequences $(\mathbf{v}_i)_{i \in \mathbb{N}} \subset \mathcal{S}$ and $(t_i)_{i \in \mathbb{N}} \subset (0, \infty)$.

Let us consider sequences $(\mathbf{v}_i)_{i \in \mathbb{N}}$ corresponding to series of the form⁵

$$\mathbf{v}_i = \sum_{j=0}^{\infty} t_i^j \mathbf{w}_j,$$

for some pairs of polynomials $\mathbf{w}_j \in \mathbb{R}[x]_d^2$ that remain to be chosen. In order for $\mathbf{v}_i$ to converge to $\mathbf{u}$, we have to take a sequence $(t_i)_{i \in \mathbb{N}}$ converging to 0 and $\mathbf{w}_0 = \mathbf{u}$. Following that, if we want to show that a vector $\mathbf{v} \in \mathbb{R}[x]_d^2$ is tangent to $\mathcal{S}$ at $\mathbf{u}$, $\mathbf{w}_1$ needs to be chosen as $\mathbf{v}$ to ensure the convergence of $\left(\frac{\mathbf{v}_i - \mathbf{u}}{t_i}\right)$ to $\mathbf{v}$.

It remains to check that $\mathbf{v}_i \in \mathcal{S}$ for all $i \in \mathbb{N}$. For this, we must have (after replacing $\mathbf{w}_0$ by $\mathbf{u}$ and $\mathbf{w}_1$ by $\mathbf{v}$)

$$\begin{aligned} \mathbf{b} &= \mathcal{L}_2(\sigma(\mathbf{v}_i)) \\ &= \mathcal{L}_2 \left(\sigma(\mathbf{u}) + 2t_i \mathcal{A}_{\mathbf{u}}(\mathbf{v}) + t_i^2 \sigma(\mathbf{v}) + 2 \sum_{j=2}^{\infty} t_i^j \mathcal{A}_{\mathbf{u}}(\mathbf{w}_j) + 2t_i \sum_{j=2}^{\infty} t_i^j \mathcal{A}_{\mathbf{v}}(\mathbf{w}_j) + \sum_{j=2}^{\infty} t_i^{2j} \sigma(\mathbf{w}_j) \right). \end{aligned}$$

Recall that $\mathbf{u} \in \mathcal{S}$, and therefore $\mathcal{L}_2(\sigma(\mathbf{u})) = \mathbf{b}$. Hence, the membership condition for $\mathbf{v}_i$ to $\mathcal{S}$ can be expressed as:

$$\mathbf{0} = \mathcal{L}_2 \left(2\mathcal{A}_{\mathbf{u}}(\mathbf{v}) + t_i \sigma(\mathbf{v}) + 2 \sum_{j=2}^{\infty} t_i^{j-1} \mathcal{A}_{\mathbf{u}}(\mathbf{w}_j) + 2 \sum_{j=2}^{\infty} t_i^j \mathcal{A}_{\mathbf{v}}(\mathbf{w}_j) + \sum_{j=2}^{\infty} t_i^{2j-1} \sigma(\mathbf{w}_j) \right), \quad (5.3)$$

for all $i \in \mathbb{N}$. The mapping $t \rightarrow f(t) := \mathcal{L}_2(2\mathcal{A}_{\mathbf{u}}(\mathbf{v}) + t\sigma(\mathbf{v}) + \dots)$ being continuous, the value of $f(\lim_{i \rightarrow \infty} t_i)$ is equal to $\lim_{i \rightarrow \infty} f(t_i)$ [Tes19, Lemma B.10]. Using (5.3),

$$\mathcal{L}_2(\mathcal{A}_{\mathbf{u}}(\mathbf{v})) = \frac{1}{2} f \left(\lim_{i \rightarrow \infty} t_i \right) = \frac{1}{2} \lim_{i \rightarrow \infty} f(t_i) = \lim_{i \rightarrow \infty} \mathbf{0} = \mathbf{0}.$$

⁵Note that not all sequence can be expressed in this way. As of today, we are not sure the conditions (5.4) describe entirely the tangent cone. Yet, if it is satisfied, we are sure to have found a tangent vector.

Removing it from the expression (5.3) and dividing by t_i , we get

$$\mathbf{0} = \mathcal{L}_2 \left(\sigma(\mathbf{v}) + 2 \sum_{j=2}^{\infty} t_i^{j-2} \mathcal{A}_u(\mathbf{w}_j) + 2t_i \sum_{j=2}^{\infty} t_i^{j-2} \mathcal{A}_v(\mathbf{w}_j) + \sum_{j=2}^{\infty} t_i^{2j-2} \sigma(\mathbf{w}_j) \right).$$

Applying the same idea recursively, we obtain the following constraints on $\mathbf{v}$ and $\mathbf{w}_j$:

$$\mathcal{L}_2(\mathcal{A}_u(\mathbf{v})) = \mathbf{0}, \quad (5.4a)$$

$$\mathcal{L}_2(\sigma(\mathbf{v}) + 2\mathcal{A}_u(\mathbf{w}_2)) = \mathbf{0}, \quad (5.4b)$$

$$\mathcal{L}_2(\mathcal{A}_u(\mathbf{w}_j) + \mathcal{A}_v(\mathbf{w}_{j-1})) = \mathbf{0}, \quad \forall j \geq 3, j \text{ odd}, \quad (5.4c)$$

$$\mathcal{L}_2(2\mathcal{A}_u(\mathbf{w}_j) + 2\mathcal{A}_v(\mathbf{w}_{j-1}) + \sigma(\mathbf{w}_{j/2})) = \mathbf{0}, \quad \forall j \geq 4, j \text{ even}. \quad (5.4d)$$

If we find $\mathbf{v} \in \mathbb{R}[x]_d^2$ and $\mathbf{w}_j \in \mathbb{R}[x]_d^2$, $j = 2, 3, \dots$, satisfying those conditions, then $\mathbf{v}$ is a tangent vector of $\mathcal{S}$ at $\mathbf{u}$.

5.2.2 Comments on the tangent cone

Several comments can be made regarding the constraints (5.4).

- A. First, (5.4a) implies that $\mathcal{A}_u(\mathbf{v})$ has to belong to the kernel of $\mathcal{L}_2$. In fact, this corresponds exactly to asking for $\mathbf{v}$ to belong to the kernel of the candidate local defining function h we analyzed in Section 5.1;
- B. Next, if $\sigma(\mathbf{v})$ belongs to the kernel of $\mathcal{L}_2$, then choosing $\mathbf{w}_j$ as the zero polynomial pair for all $j \geq 2$ satisfies equations (5.4b)–(5.4d);
- C. Third, if $\text{Im } \mathcal{A}_u \subseteq \ker \mathcal{L}_2$, the condition (5.4a) is automatically satisfied, and condition (5.4b) becomes $\mathcal{L}_2(\sigma(\mathbf{v})) = \mathbf{0}$. For the conditions (5.4c)–(5.4d), we may again choose $\mathbf{w}_j$ as the zero polynomials pair for all $j = 2, 3, \dots$;
- D. On the other hand, if $\ker \mathcal{L}_2 \subseteq \text{Im } \mathcal{A}_u$, condition (5.4b) implies that $\sigma(\mathbf{v})$ must belong to the image of $\mathcal{A}_u$. For this, the polynomials of $\mathbf{v}$ must share the same common roots as the ones of $\mathbf{u}$ (including their multiplicity). If $\mathbf{v}$ satisfies this constraint, since $\mathcal{A}_u$ is onto its image, there exists pairs of polynomials $\mathbf{w}_j$, $j \geq 2$, satisfying equations (5.4b)–(5.4d). Yet, $\mathbf{v}$ still has to satisfy (5.4a);
- E. Note that, if $\text{Im } \mathcal{A}_u = \ker \mathcal{L}_2$, the two points hereabove correspond: $\sigma(\mathbf{v})$ must belong to the kernel of $\mathcal{L}_2$, which is equal to $\text{Im } \mathcal{A}_u$;
- F. $\mathcal{L}_2$ and $\mathcal{A}_u$ being linear applications, if their composition is surjective, for all $\mathbf{v} \in \ker \mathcal{L}_2 \circ \mathcal{A}_u$, there exist polynomials $\mathbf{w}_j$ satisfying the equations (5.4). For this end, first solve the equation (5.4b) (which has a solution since $\mathcal{L}_2 \circ \mathcal{A}_u$ is surjective), then iteratively solve the equations (5.4c) and (5.4d) to obtain the next odd and even $\mathbf{w}_j$ respectively.
- G. Finally, if the kernel of $\mathcal{L}_2$ contains only the zero polynomial (i.e., $\mathcal{L}_2$ is bijective), $\mathbf{v}$ has to be chosen in the kernel of $\mathcal{A}_u$ (due to (5.4a)). Moreover, there should exist a pair of polynomials $\mathbf{w}_2$ such that $2\mathcal{A}_u(\mathbf{w}_2) = -\sigma(\mathbf{v})$ (this corresponds to

condition (5.4b)). This implies that $\sigma(\mathbf{v})$ has to belong to the image of $\mathcal{A}_{\mathbf{u}}$, or in other words that $\sigma(\mathbf{v})$ vanishes at all common roots of $\mathbf{u}$. For this to happen, both polynomials of $\mathbf{v}$ must vanish at those common roots (as well as their derivatives if some roots are of multiplicity higher than one). The constraints (5.4c)–(5.4d) for finding the pairs of polynomials $\mathbf{w}_j$, $j \geq 3$ will then admit solutions. This is a particular case of item D.

5.2.3 Tangent cone of smooth embedded submanifolds

These observations are interesting, as they allow us to verify a fundamental result on the embedded submanifolds of Chapter 4, which is that if $\mathcal{S}$ is a smooth embedded submanifold locally around some point, then its tangent cone coincides with the *tangent space* at that point [RW98, Example 6.8]. The *tangent space* to a manifold at a point is a linear approximation of the manifold around the point. Geometrically, one can think of the tangent line to a curve or to the tangent space to a sphere at some point. The tangent space will be used in Chapter 6 to define a Riemannian gradient, as it requires the notion of scalar product on a linear space. When we consider embedded submanifolds of Euclidean spaces, a simple definition of the tangent space is given by Definition 5.8. More information about the origin of this definition can be found in [Bou23, Section 3.2].

Definition 5.8 – Tangent space

Let $\mathcal{M}$ be an embedded submanifold of a vector space $\mathcal{E}$, and let $x \in \mathcal{M}$. If $\mathcal{M}$ is an open submanifold, then the *tangent space* of $\mathcal{M}$ at x , denoted $\mathcal{T}_x\mathcal{M}$, is identified as the vector space $\mathcal{E}$. Otherwise, $\mathcal{T}_x\mathcal{M}$ is equal to the kernel of $Dh(x)$ for some local defining function h of $\mathcal{M}$ at x .

Remark 5.9. An interesting property about the tangent space to a manifold at a point is that it is a vector subspace of $\mathcal{E}$, of the same dimension as the manifold [Bou23, Def 3.16]. $\triangle$

Coprime polynomials (Section 4.3) Consider the case of a pair of coprime polynomials $\mathbf{u}$. In Section 4.3, we showed that the set $\mathcal{S}$ is locally a smooth manifold by using the local defining function h (5.1). Since $\mathbf{u}$ is coprime, $\mathcal{A}_{\mathbf{u}}$ is surjective, implying the surjectivity of $\mathcal{L}_2 \circ \mathcal{A}_{\mathbf{u}}$. Hence, by item F., the tangent cone at $\mathbf{u}$ is the kernel of $\mathcal{L}_2 \circ \mathcal{A}_{\mathbf{u}}$, which corresponds to the tangent space given by the kernel of $Dh(\mathbf{u})$ (Lemma 5.1) according to Definition 5.8.

Bijective operator $\mathcal{L}_2$ (Section 4.2) Next, consider a bijective application $\mathcal{L}_2$. In this case, we can remove the operator $\mathcal{L}_2$ from the four conditions (5.4) (e.g., the first one reduces to $\mathbf{v} \in \ker \mathcal{A}_{\mathbf{u}}$). If $\mathbf{u}$ is coprime, $\mathcal{A}_{\mathbf{u}}$ is surjective and conditions (5.4b)–(5.4d) can be satisfied for all vector $\mathbf{v} \in \mathbb{R}[x]_d^2$. Therefore, the tangent cone is fully described by (5.4a): a vector $\mathbf{v} \in \mathbb{R}[x]_d^2$ is a tangent vector if and only if it lies in the kernel of $\mathcal{A}_{\mathbf{u}}$, which is the one-dimensional vector space corresponding to the kernel of the local defining function h (5.1).

If $\mathbf{u}$ is not coprime, $\mathbf{v}$ still needs to belong to the kernel of $\mathcal{A}_{\mathbf{u}}$, which is now a $K + 1$ dimensional vector space (where K is the number of common roots). Yet, to satisfy condition (5.4b), $\sigma(\mathbf{v})$ must belong to the image of $\mathcal{A}_{\mathbf{u}}$. For this, the polynomials of $\mathbf{v}$ must

vanish⁶ at the K common roots of the polynomials of $\mathbf{u}$. These K linear (independent) constraints reduce the set of allowed tangent vectors $\mathbf{v}$ to a one-dimensional vector space. Once again, this corresponds to the tangent space (recall that the set $\mathcal{S}$ is a union of disjoint circles, for which the tangent space at a point can be seen as the tangent line to the one-dimensional circle). Note that, if the tangent vectors $\mathbf{v}$ satisfy these conditions, there always exist pairs of polynomials $\mathbf{w}_j$, $j \geq 3$ such that the constraints (5.4c)–(5.4d) are satisfied.

Polynomials with one common root (Example 5.2) Finally, in Example 5.2, we found a case in which the function h is not a local defining function, yet we subsequently showed in Remark 5.3 that this case never happens. Specifically, we analyzed what happens if $\ker \mathcal{L}_2 \subseteq \text{Im } \mathcal{A}_u$ at some point $\mathbf{u} \in \mathcal{S}$. In that case, following item D., both polynomials of the tangent vector $\mathbf{v}$ must share the same common roots (including their multiplicity) as the ones of $\mathbf{u}$, and $\mathbf{v}$ must belong to the kernel of $\mathcal{L}_2 \circ \mathcal{A}_u$. Those constraints being linear on $\mathbf{v}$, the tangent cone is linear as expected.

5.2.4 Tangent cone of our counterexample

However, problems arise when the image of $\mathcal{A}_u$ is strictly contained in the kernel of $\mathcal{L}_2$: $\text{Im } \mathcal{A}_u \subset \ker \mathcal{L}_2$. When this occurs, item C. indicates that $\sigma(\mathbf{v})$ has to belong to the kernel of $\mathcal{L}_2$. This can be expressed as:

$$\sigma(\mathbf{v}) = \mathcal{A}_u(\mathbf{w}) + \sum_{k=1}^{K-m} \alpha_k \epsilon_k,$$

for some pair of polynomials $\mathbf{w} \in \mathbb{R}[x]_d^2$ and coefficients $\alpha_k \in \mathbb{R}$. Here, $\epsilon_k \in \mathbb{R}[x]_{2d}$ are the $K - m$ polynomials⁷ in the kernel of $\mathcal{L}_2$ that do not belong to the image of $\mathcal{A}_u$, with K the number of common roots between the polynomials of $\mathbf{u}$. Note that for $K - m$ to be positive, we must have $K > m \geq 1$, i.e., K is at least 2. We can rewrite $\mathcal{A}_u(\mathbf{w})$ as $g \cdot \tilde{w}$, where $g \in \mathbb{R}[x]_K$ contains the common roots of $\mathbf{u}$, and $\tilde{w}$ is a polynomial in $\mathbb{R}[x]_{2d-K}$.

Therefore, we are searching for a pair of polynomials $\mathbf{v} \in \mathbb{R}[x]_d^2$ such that there exist $w \in \mathbb{R}[x]_{2d-K}$ and $\alpha_k \in \mathbb{R}$ for which the equality $\sigma(\mathbf{v}) - \sum_{k=1}^{K-m} \alpha_k \epsilon_k = g \cdot \tilde{w}$ is satisfied. This is strictly equivalent to asking that the polynomial on the left hand side vanishes at the roots of g , which gives us K conditions⁸:

$$v_1(x_i)^2 + v_2(x_i)^2 = \sum_{k=1}^{K-m} \alpha_k \epsilon_k(x_i), \quad i = 1, \dots, K.$$

The set of vectors $\mathbf{v}$ satisfying these conditions is not a vector space. Indeed, assume $\mathbf{v}$ and $\mathbf{v}'$ satisfy these conditions; then, there may not exist $\tilde{\alpha}_k$ such that, for all $i = 1, \dots, K$,

⁶If there is a root of multiplicity $\mu > 1$, the $\mu - 1$ first derivatives of the polynomial must also vanish at the root. Furthermore, these constraints on the values of the derivatives are linearly independent from the constraints on the values of the polynomial. This ensures that we have K linearly independent constraints.

⁷The kernel of $\mathcal{L}_2$ is of dimension $2d + 1 - m$, and the one of the image of $\mathcal{A}_u$ is $2d + 1 - K$, where K is the number of common roots between the polynomials of $\mathbf{u}$ (Lemma 4.1)

⁸For simplicity, we assume the roots to be of multiplicity one, but this can be generalized to roots of higher multiplicity.

we have

$$(v_1(x_i) + v'_1(x_i))^2 + (v_2(x_i) + v'_2(x_i))^2 = \sum_{k=1}^{K-m} \tilde{\alpha}_k \epsilon_k(x_i).$$

Remark 5.10. Following Remark 5.4, the constraint vector $\mathbf{b}$ is the zero vector. The tangent cone to $\mathcal{S}$ at $\mathbf{u}$ being given by the set of pairs of polynomials $\mathbf{v} \in \mathbb{R}[x]_d^2$ such that $\sigma(\mathbf{v})$ belongs to the kernel of $\mathcal{L}_2$, it corresponds exactly to the set $\mathcal{S}$ itself! $\triangle$

Example 5.11. Let us continue Example 5.5. Recall that the operator $\mathcal{L}_2$ we considered (which takes a 4-degree polynomial and returns the difference between its evaluation at 3 and at 1) was such that the image of $\mathcal{A}_v$ is contained in the kernel of $\mathcal{L}_2$, for $\mathbf{v} = ((x-1)(x-3), 2(x-1)(x-3))$. Hence, according to Remark 5.10, the tangent cone to $\mathcal{S}$ at $\mathbf{v}$ is the whole set $\mathcal{S}$.

However, the set $\mathcal{S}$ is not a linear space. Indeed, the set $\mathcal{S}$ contains all the pairs of polynomials such that their sum of squares evaluated at 1 is equal to its evaluation at 3. In particular, if we consider the pair of polynomials $\mathbf{u} = (u_1, u_2) = (x(x-2), -x^2 + 3x + 1)$ and the pair $(-u_1, u_2)$, they both belong to $\mathcal{S}$; but their sum, which is equal to $(0, 2u_2)$ does not.

For a visual interpretation, let us consider the set

$$\tilde{\mathcal{S}} = \left\{ \alpha R(\theta) \mathbf{u} + (x-1)(x-3) \begin{pmatrix} \beta \\ \gamma \end{pmatrix} \mid \alpha, \beta, \gamma \in \mathbb{R}, \theta \in [0, 2\pi[\right\},$$

where $R(\cdot)$ is a rotation matrix. This is a subset of $\mathcal{S}$, as

$$\begin{aligned} & \sigma \left(\alpha R(\theta) \mathbf{u} + (x-1)(x-3) \begin{pmatrix} \beta \\ \gamma \end{pmatrix} \right) \\ &= \underbrace{\alpha^2 \sigma(\mathbf{u})}_{=\tilde{\alpha} \sigma(\mathbf{u})} + \underbrace{2\alpha(x-1)(x-3) \mathbf{u}^T R(\theta)^T \begin{pmatrix} \beta \\ \gamma \end{pmatrix}}_{=(x-1)(x-3)p \in \text{Im } \mathcal{A}_v} + (x-1)^2(x-3)^2(\beta^2 + \gamma^2). \end{aligned}$$

Being the sum of two elements in the kernel of $\mathcal{L}_2$, it lies in the kernel of $\mathcal{L}_2$.

The set $\tilde{\mathcal{S}} \subseteq \mathcal{S}$ is represented on Figure 5.1. It corresponds to a vector space (the one given by $(x-1)(x-3) \begin{pmatrix} \alpha \\ \beta \end{pmatrix}$), from which a double cone emerges at every point (corresponding to $\alpha R(\theta) \mathbf{u}$). $\diamond$




Figure 5.1: Subset $\tilde{\mathcal{S}}$ of the tangent cone to the set $\mathcal{S}$ described in Example 5.5 at the point $\mathbf{v} = ((x-1)(x-3), 2(x-1)(x-3))$. It is important to keep in mind that this is a 3-dimensional representation of a subset of a 6-dimensional space $(\mathbb{R}[x]_2^2)$. The line may be seen as a two-dimensional vector space, corresponding to $(x-1)(x-3) \begin{pmatrix} \beta \\ \gamma \end{pmatrix}$ for some $\beta, \gamma \in \mathbb{R}$. At each point of that space, we have a double infinite cone corresponding to $\alpha R(\theta)\mathbf{u}$, for some $\alpha \in \mathbb{R}$ and $\theta \in [0, 2\pi[$. Note that all those cones are tangent to $\mathcal{S}$ at $\mathbf{v}$, not only the one starting at $\mathbf{v}$. The set $\mathcal{S}$ is similar, but with a cone emerging in each direction of pairs of polynomials for which the sum of their squares has the same value at 1 and at 3.

5.3 A stratified structure

A closely related field to optimization on manifolds is optimization on *stratified sets*.

Definition 5.12 – Stratified set

Let $\mathcal{M}$ be a smooth manifold, and $\mathcal{S}$ be a subset of $\mathcal{M}$. The set $\mathcal{S}$ is called a *stratified set* if there exists a cover of $\mathcal{S}$ by pairwise disjoint smooth submanifolds of $\mathcal{M}$, which lie in $\mathcal{S}$, called *strata*. Furthermore, $\mathcal{S}$ has to satisfy the *condition of the frontier*: for each stratum $\mathcal{X}$ of $\mathcal{S}$, its frontier $(\bar{\mathcal{X}} - \mathcal{X}) \cap \mathcal{S}$ is a union of strata.

In other words, a stratified set is a union of smooth manifolds that are nested inside one another, in the sense that the frontier of a stratum is a union of strata. It is quite surprising that, while Riemannian optimization on smooth manifolds has been well studied for a long time, the first algorithms on stratified sets were only discovered a few years ago by G. Olikier et al. [OGA23] and N. Boumal et al. [LKB23].

A major issue when optimizing on stratified sets is the emergence of *apocalypses*, from the Greek ἀποκάλυψις, meaning the act of discovering, or a revelation [Bai89]. An *apocalypse* appears when, for some function $f : \mathcal{S} \rightarrow \mathbb{R}$ (for some set $\mathcal{S}$), the stationarity measure along iterates $(x_k)_{k \in \mathbb{N}}$ (obtained via some optimization method) converges to zero,

but the iterates converge to a point which is not stationary⁹ [LKB23, Def. 2.7]:

$$\lim_{k \rightarrow \infty} \text{grad } f(x_k) = 0, \quad \text{grad } f \left(\lim_{k \rightarrow \infty} x_k \right) \neq 0.$$

So, apocalypses live up to their name: when reaching the limit of the iterates, a new research direction is revealed, while it was unknown until being reached. The points corresponding to the limit of those iterates are called *apocalyptic points*.

5.3.1 Stratification on our example

Apocalyptic points are therefore a geometric property of the constraint sets. Boumal et al. give the intuition behind apocalyptic points as accumulation points of some sequence, for which the limit of the tangent cone does not contain the tangent cone at the limit. This is the case for several stratified sets, like the one of the bounded-rank matrices [OGA23; LKB23]. In both papers, the authors propose algorithms tackling the issue of apocalyptic points, allowing the convergence to a true stationary point.

In what follows, we show that the set described in Examples 5.5 and 5.11 is a stratified set.

Example 5.13. Once again, consider the Examples 5.5 and 5.11. Recall that the kernel of $\mathcal{L}_2$ is given by the set of polynomials $p \in \mathbb{R}[x]_4$ such that $p(1) = p(3)$. We can express $\mathcal{S}$ as the union of an infinite number of disjoint sets,

$$\begin{aligned} \mathcal{S}_1 &= \left\{ \mathbf{u} \in \mathbb{R}[x]_2^{2,c} \mid \mathcal{L}_2(\sigma(\mathbf{u})) = 0 \right\}; \\ \mathcal{S}_2^\alpha &= (x - \alpha) \left\{ \mathbf{u} \in \mathbb{R}[x]_1^{2,c} \mid \mathcal{L}_2((x - \alpha)^2 \sigma(\mathbf{u})) = 0 \right\}, \quad \alpha \in \mathbb{R}; \\ \mathcal{S}_3^{\alpha,\beta} &= (x - \beta)(x - \alpha) \left\{ \mathbf{u} \in \mathbb{R}[x]_0^2 \mid \mathcal{L}_2((x - \alpha)^2 (x - \beta)^2 \sigma(\mathbf{u})) = 0 \right\}, \quad \alpha \leq \beta \in \mathbb{R}. \end{aligned}$$

Indeed, we can show that there exists an infinite number of sets $\mathcal{S}_2^\alpha$ that are non empty. Let $\alpha \in \mathbb{R} \setminus \{1, 3\}$ (see below why we remove 1 and 3), and fix the values¹⁰ of $u_1(1), u_1(3), u_2(1)$ and $u_2(3)$ such that

$$\begin{aligned} u_1(1)(1 - \alpha) &= u_2(3)(3 - \alpha) = \zeta, \\ u_1(3)(3 - \alpha) &= u_2(1)(1 - \alpha) = \xi, \end{aligned}$$

with $\zeta \neq \xi \neq 0$. By construction, $\mathbf{u}$ belongs to $\mathcal{S}_2^\alpha$ since

$$(1 - \alpha)^2 (u_1^2(1) + u_2^2(1)) = (3 - \alpha)^2 (u_1^2(3) + u_2^2(3)) = \zeta^2 + \xi^2.$$

Since u_1 and u_2 are degree-1 polynomials, they are fully described by these relations, and they are coprime since $\zeta \neq \xi$.

Note that if $\alpha = 1$ (resp. 3), we must have $u_2(3) = u_1(3) = 0$ (resp. $u_1(1) = u_2(1) = 0$), implying that $\mathcal{S}_2^1$ and $\mathcal{S}_2^3$ are empty.

⁹Here, the gradient corresponds to the projection of the Euclidean gradient onto the tangent cone (Definition 5.7), and not onto the tangent space (Definition 5.8) (which was the case of the Riemannian gradient — Definition 6.2)

¹⁰The pair of polynomials $\mathbf{u}$ is now a general element of $\mathbb{R}[x]_1^{2,c}$, and no more the pair of polynomials we considered in Examples 5.5 and 5.11.

Another representation of $\mathcal{S}$ is as the union of exactly two disjoint manifolds:

$$\begin{aligned}\mathcal{M}_1 &= \mathcal{S}_1 \cup \left(\bigcup_{\alpha \in \mathbb{R} \setminus \{1,3\}} \mathcal{S}_2^\alpha \right) \cup \left(\bigcup_{(\alpha,\beta) \in \mathbb{R}^2 \setminus \{(1,3)\}} \mathcal{S}_3^{\alpha,\beta} \right), \\ \mathcal{M}_2 &= \mathcal{S}_3^{1,3}.\end{aligned}$$

It is easily seen that $\mathcal{M}_2 = \mathcal{S}_3^{1,3} = (x-1)(x-3)\mathbb{R}[x]_0^2$ is a smooth manifold of dimension 2, according to Remark 4.10.

For $\mathcal{M}_1$, this requires a bit more work. Consider our usual candidate local defining function

$$h : \mathbb{R}[x]_d^2 \setminus (x-1)(x-3)\mathbb{R}[x]_0^2 \rightarrow \mathbb{R} : \mathbf{u} \rightarrow \mathcal{L}_2(\sigma(\mathbf{u})).$$

Then, $h^{-1}(0) = \mathcal{M}_1$. In Chapter 4, we studied the rank of its differential Dh at several points:

- At a coprime point $\mathbf{u} \in \mathcal{S}_1$, Theorem 4.9 tells us that the rank of $Dh(\mathbf{u})$ is $m = 1$;
- At a point $\mathbf{u} \in \mathcal{S}_2^\alpha$ (for some $\alpha \notin \{1,3\}$) where the two polynomials share exactly one root, we saw in Section 5.1.1 that the only way for the rank of $Dh(\mathbf{u})$ to be different from m is that all the polynomials in the kernel of $\mathcal{L}_2$ vanish at that common root, which is not the case;
- Finally, at a point $\mathbf{u} \in \mathcal{S}_3^{\alpha,\beta}$ (for some $\alpha, \beta \notin \{1,3\}$), a similar reasoning in Section 5.1.2 showed that the rank of $Dh(\mathbf{u})$ is different from m only if $\text{Im } \mathcal{A}_\mathbf{u}$ is contained in the kernel of $\mathcal{L}_2$, which is not possible since it is already the case of $\text{Im } \mathcal{A}_\mathbf{v}$, for $\mathbf{v} = ((x-1)(x-3), 2(x-1)(x-3))$, and α, β are different from 1 and 3.

Hence, at all the points $\mathbf{u}$ in $\mathcal{M}_1$, the differential $Dh(\mathbf{u})$ is full rank, implying that $\mathcal{M}_1$ is a smooth manifold of dimension $2d + 1 = 5$.

Furthermore, $\mathcal{S}$ satisfies the condition of the frontier:

- $\mathcal{M}_2$ has an empty frontier as it is a closed set;
- The frontier of $\mathcal{M}_1$ is the stratum $\mathcal{M}_2$. Indeed, for all the points in $\mathcal{M}_2$, we saw in Example 5.11 that the tangent cone to $\mathcal{S}$ is the whole set $\mathcal{S}$ at these points. If we consider $\mathbf{u} \in \mathcal{M}_2$, we saw that for all $\mathbf{v} \in \mathcal{S}$, the sequence $(\mathbf{u} + t_i \mathbf{v})_{i \in \mathbb{N}}$ converges to $\mathbf{u}$ as $t_i \rightarrow 0$. In particular, if $\mathbf{v} \in \mathcal{M}_1$, then $\mathbf{u} + t_i \mathbf{v}$ belongs to $\mathcal{M}_1$ for all $i \in \mathbb{N}$, showing that $\mathbf{u}$ is a boundary point of $\mathcal{M}_1$. Since $\mathcal{S} = \mathcal{M}_1 \cup \mathcal{M}_2$, the frontier of $\mathcal{M}_1$ is fully described by $\mathcal{M}_2$.

According to Definition 5.12, $\mathcal{S}$ is therefore a stratified set with two strata. ◇

5.3.2 A remedy when $\mathbf{b} = 0$

A natural question now is whether the set described in Examples 5.5, 5.11 and 5.13 has apocalyptic points. Yet, a simple observation allows us to sidestep this question for the moment. The example was constructed so that $\text{Im } \mathcal{A}_\mathbf{v} \subset \ker \mathcal{L}_2$ for some pair of polynomials $\mathbf{v} \in \mathcal{S}$. As it has already been said, this implies that the constraint vector $\mathbf{b}$

must be the zero vector. Hence, if a pair of polynomials $\mathbf{u}$ belongs to the feasible set $\mathcal{S}$, all the multiples of $\mathbf{u}$ belong to $\mathcal{S}$:

$$\mathcal{L}_2(\sigma(\mathbf{u})) = \mathbf{0} \implies \mathcal{L}_2(\sigma(\lambda\mathbf{u})) = \lambda^2 \mathcal{L}_2(\sigma(\mathbf{u})) = \mathbf{0} \quad \forall \lambda \in \mathbb{R}.$$

On the objective (which we seek to minimize), this can be reflected in two different ways:

- If $\mathcal{L}_1(\sigma(\mathbf{u})) > 0$, then $\mathbf{u}$ is not optimal as the zero pair of polynomials reach a zero value ($\mathcal{L}_1((0, 0)) = 0$);
- If $\mathcal{L}_1(\sigma(\mathbf{u})) < 0$, then the problem is unbounded, as multiplying $\mathbf{u}$ by some factor λ increasingly larger leads to an increasingly smaller value.

Hence, either the zero pair of polynomials is optimal, or the problem is unbounded. This is valid whenever the constraint vector $\mathbf{b}$ is the zero vector. Therefore, there is no need to apply an optimization algorithm, and there is no need to worry about potential apocalypses for now.

Chapter 6

First order algorithm for sum-of-squares problem

After this long journey in the world of polynomials, we may now go back to the original problem we considered, which was the optimization problem (1.2). All the analysis of the feasible set realized in the previous chapters had as objective to rely on a Riemannian optimization algorithm. The most classical one is the Riemannian Gradient Descent [Bou23, Chapter 4].

Let $\mathcal{M}$ be a smooth manifold, and $f : \mathcal{M} \rightarrow \mathbb{R}$ be a smooth function¹ on $\mathcal{M}$. Then, the Riemannian Gradient Descent method² is given by [Bou23, Algorithm 4.1]

Input: An initial point x_0 on the manifold $\mathcal{M}$
for $k \leftarrow 1, 2, \dots$ **do**
 Pick a stepsize $\alpha_k > 0$;
 $x_{k+1} = R_{x_k}(-\alpha_k \text{grad } f(x_k))$.
end

Algorithm 1: Riemannian Gradient Descent

In comparison with the standard Gradient Descent algorithm on a Euclidean space, there are two new concepts that need explanation. The first one is the notion of *Riemannian gradient* $\text{grad } f(x)$ (at a point $x \in \mathcal{M}$). This is a generalization of the Euclidean gradient lying on the tangent space to the manifold at x . We give more information about the Riemannian gradient in Section 6.1.

The second one is the notion of *retraction*. Indeed, a key property of an optimization algorithm on manifolds is that all the iterates lie on the manifold. Yet, the Gradient Descent would produce an iterate $x_{k+1} = x_k - \alpha_k \text{grad } f(x_k)$ lying on the tangent space (Definition 5.8) to $\mathcal{M}$ at x_k . Hence, we need an operator $R_{x_k} : \mathcal{T}_{x_k} \mathcal{M} \rightarrow \mathcal{M}$ that maps a point on the tangent space to the manifold, while keeping similar information as the initial

¹If $\mathcal{M}$ is an embedded submanifold of a linear space $\mathcal{E}$, a function $f : \mathcal{M} \rightarrow \mathbb{R}$ is *smooth* if and only if it is the restriction to $\mathcal{M}$ of some smooth function $\tilde{f} : \mathcal{U} \rightarrow \mathbb{R}$, where $\mathcal{U}$ is a neighbourhood of $\mathcal{M}$ in $\mathcal{E}$ [Bou23, Prop 3.31]. Without needing to understand fully this definition, we are concerned in this thesis by the function $\mathbb{R}[x]_d^2 \rightarrow \mathbb{R} : \mathbf{u} \mapsto \mathcal{L}_1(\sigma(\mathbf{u}))$, which is smooth on $\mathbb{R}[x]_d^2$ (as a linear map and the sum-of-squares operator are smooth and smoothness is preserved by composition). Hence, its restriction to any embedded submanifold of $\mathbb{R}[x]_d^2$ is smooth.

²In what follows, we will not detail the choice of the stepsize α_k . A discussion on the different possible choices is given in [Bou23, Section 4.3].

point (e.g., the gradient is usually seen as the direction of the steepest ascent — hence, it may not be mapped into a direction completely opposite, such as the steepest descent). More detail about the retraction are given in Section 6.2.

An implementation of the Riemannian Gradient Descent for problem (1.2) (relying on the `Manopt.jl` package) with the Riemannian gradient and the retraction presented in the next sections is available at <https://github.com/CharlesVanHees/Manifold-SOS>.

6.1 Riemannian gradient

6.1.1 Definition

In Euclidean spaces, a classical interpretation of the gradient is that it corresponds to the direction of the steepest ascent. This follows from the definition of the gradient:

Definition 6.1 – Gradient

Let $f : \mathcal{E} \rightarrow \mathbb{R}$ be a smooth function defined on a Euclidean space $\mathcal{E}$. The *gradient* of f at a point $x \in \mathcal{E}$ is the unique vector $\nabla f(x) \in \mathcal{E}$ such that, for all $v \in \mathcal{E}$,

$$Df(x)[v] = \langle \nabla f(x), v \rangle,$$

where $Df(x)$ is the differential operator of f at x , defined by

$$Df(x) : \mathcal{E} \rightarrow \mathbb{R} : v \mapsto \lim_{t \rightarrow 0} \frac{f(x + tv) - f(x)}{t}.$$

According to this definition, the differential of f at x , which represents an infinitesimal variation in a given direction, is maximized in a direction parallel to the gradient of f at x , according to the Cauchy-Schwarz inequality.

For a smooth manifold $\mathcal{M}$, there are two issues with that definition. First, the gradient is intrinsically linked with an inner product, defined on a linear space. To tackle this issue, we rely on the fact that the tangent space to a manifold at a point is a linear space. If we define for all $x \in \mathcal{M}$ an inner product $\langle \cdot, \cdot \rangle_x$ on $\mathcal{T}_x \mathcal{M}$, that varies smoothly with x (in a sense that we do not develop here — see [Bou23, Section 3.7]), $\mathcal{M}$ is said to be a *Riemannian manifold*. On a Riemannian manifold, we can define the notion of *Riemannian gradient* [Bou23, Def 3.58]:

Definition 6.2 – Riemannian gradient

Let $f : \mathcal{M} \rightarrow \mathbb{R}$ be smooth on a Riemannian manifold $\mathcal{M}$. The *Riemannian gradient* of f at a point $x \in \mathcal{M}$ is the unique tangent vector $\text{grad } f(x) \in \mathcal{T}_x \mathcal{M}$ such that, for all $v \in \mathcal{T}_x \mathcal{M}$,

$$Df(x)[v] = \langle \text{grad } f(x), v \rangle_x,$$

where $\langle \cdot, \cdot \rangle_x$ is the Riemannian metric.

Yet, the differential of f at x in a direction v is not well defined, as f may not be defined at $x + tv$ if it does not belong to $\mathcal{M}$. Note that every vector $v \in \mathcal{E}$ can be decomposed into a vector in the tangent space to $\mathcal{M}$ at x , and a normal component. As the variation of a function is null in a normal decomposition, the differential of a function at a point is

defined on the tangent space to the manifold at that point:

$$Df(x) : \mathcal{T}_x\mathcal{M} \rightarrow \mathbb{R}.$$

A precise definition of the differential of a smooth function at a point is given in [Bou23, Def 3.34].

For what concerns us, we may rely on an interesting property. For an embedded submanifold $\mathcal{M}$ of a Euclidean space $\mathcal{E}$, the tangent spaces $\mathcal{T}_x\mathcal{M}$ are linear subspaces of $\mathcal{E}$, for all $x \in \mathcal{M}$. Hence, a natural choice is to take as inner product on $\mathcal{T}_x\mathcal{M}$ a restriction of the inner product on $\mathcal{E}$. In that case, we talk about a *Riemannian submanifold*, and we have the following property about the Riemannian gradient [Bou23, Prop 3.61]:

Lemma 6.3

Let $\mathcal{M}$ be a Riemannian submanifold of a Euclidean space $\mathcal{E}$ endowed with the metric $\langle \cdot, \cdot \rangle$, and let $f : \mathcal{M} \rightarrow \mathbb{R}$ be a smooth function. The Riemannian gradient of f at $x \in \mathcal{M}$ is given by

$$\text{grad } f(x) = \text{Proj}_x \left(\nabla \bar{f}(x) \right),$$

where $\bar{f}$ is a smooth extension of f to a neighbourhood of $\mathcal{M}$ in $\mathcal{E}$, and Proj_x is the orthogonal projector onto the tangent space to $\mathcal{M}$ at x .

6.1.2 Riemannian gradient for polynomials

On the linear space of pairs of polynomials of degree d , $\mathbb{R}[x]_d^2$, a natural choice of metric is to extract the expression of the polynomials in some basis, stack the two $d+1$ -dimensional vectors into a $2d+2$ dimensional one for each pair, and then take the standard inner product in $\mathbb{R}^{2d+2}$:

$$\langle \cdot, \cdot \rangle_{\mathbb{R}[x]_d^2} : \mathbb{R}[x]_d^2 \rightarrow \mathbb{R} : (\mathbf{u}, \mathbf{v}) \mapsto u_1^T v_1 + u_2^T v_2,$$

where u_1, u_2 (resp. v_1, v_2) are the expression of the two polynomials of $\mathbf{u}$ (resp. $\mathbf{v}$) in some basis. Using this inner product on each tangent space, all the embedded submanifolds of $\mathbb{R}[x]_d^2$ are Riemannian submanifolds of $\mathbb{R}[x]_d^2$.

Back to our original problem (1.2), the objective is the function $f : \mathbb{R}[x]_d^2 \rightarrow \mathbb{R} : \mathbf{u} \mapsto \mathcal{L}_1(\sigma(\mathbf{u}))$, which is smooth on every embedded submanifold of $\mathbb{R}[x]_d^2$ as it is smooth on $\mathbb{R}[x]_d^2$. The Euclidean gradient of $\bar{f} : \mathbb{R}[x]_d^2 \rightarrow \mathbb{R} : \mathbf{u} \mapsto \mathcal{L}_1(\sigma(\mathbf{u}))$, a smooth extension of f to $\mathbb{R}[x]_d^2$, is given by

$$\nabla \bar{f}(\mathbf{u}) = 2L_1 A_{\mathbf{u}},$$

where L_1 and $A_{\mathbf{u}}$ are matrix representations of the operators $\mathcal{L}_1$ and $\mathcal{A}_{\mathbf{u}}$ in the same basis as the one chosen for the inner product.

According to Lemma 6.3, we therefore simply have to find the orthogonal projector onto the tangent space. Consider a smooth embedded submanifold $\mathcal{M}$ of $\mathbb{R}[x]_d^2$, a point $\mathbf{u} \in \mathcal{M}$, and a local defining function $h : \mathbb{R}[x]_d^2 \rightarrow \mathbb{R}^k$ for $\mathcal{M}$ around $\mathbf{u}$ [Bou23, Section

7.7]. The differential of h at $\mathbf{u}$ can be expressed as

$$\begin{aligned} Dh(\mathbf{u})[\mathbf{v}] &= \lim_{t \rightarrow 0} \frac{h(\mathbf{u} + t\mathbf{v}) - h(\mathbf{u})}{t} = \lim_{t \rightarrow 0} \frac{1}{t} \begin{pmatrix} h_1(\mathbf{u} + t\mathbf{v}) - h_1(\mathbf{u}) \\ \dots \\ h_k(\mathbf{u} + t\mathbf{v}) - h_k(\mathbf{u}) \end{pmatrix} \\ &= \begin{pmatrix} Dh_1(\mathbf{u})[\mathbf{v}] \\ \dots \\ Dh_k(\mathbf{u})[\mathbf{v}] \end{pmatrix} = \begin{pmatrix} \langle \nabla h_1(\mathbf{u}), \mathbf{v} \rangle \\ \dots \\ \langle \nabla h_k(\mathbf{u}), \mathbf{v} \rangle \end{pmatrix} \end{aligned}$$

Recall that the tangent space to $\mathcal{M}$ at $\mathbf{u}$ is defined as the kernel of $Dh(\mathbf{u})$ (Definition 5.8). We define the *normal space* to $\mathcal{M}$ at $\mathbf{u}$, denoted $\mathcal{N}_\mathbf{u}\mathcal{M}$, as the set of vectors in $\mathbb{R}[x]_d^2$ that are orthogonal (in the sense of the inner product in $\mathbb{R}[x]_d^2$) to all the vectors in the tangent space to $\mathcal{M}$ at $\mathbf{u}$. By definition of the tangent space, the normal space to $\mathcal{M}$ at $\mathbf{u}$ is given by³:

$$\mathcal{N}_\mathbf{u}\mathcal{M} = \left\{ \sum_{i=1}^k \alpha_i \text{grad } h_i(\mathbf{u}) \mid \alpha_1, \dots, \alpha_k \in \mathbb{R} \right\}.$$

Accordingly, we define the operator⁴

$$Dh(\mathbf{u})^* : \mathbb{R}^k \rightarrow \mathbb{R}[x]_d^2 : \alpha \mapsto \sum_{i=1}^k \alpha_i \text{grad } h_i(\mathbf{u}).$$

This allows us to express the orthogonal projection of a vector $\mathbf{v} \in \mathbb{R}[x]_d^2$ onto $\mathcal{N}_\mathbf{u}\mathcal{M}$ as the vector $\mathbf{w} \in \mathcal{N}_\mathbf{u}\mathcal{M}$ that minimizes $\|\mathbf{v} - \mathbf{w}\|^2$. Since $\mathbf{w} \in \mathcal{N}_\mathbf{u}\mathcal{M}$, there exists $\alpha^* \in \mathbb{R}^k$ such that

$$\mathbf{w} = Dh(\mathbf{u})^*[\alpha^*].$$

The vector α^* is such that

$$\alpha^* = \arg \min_{\alpha \in \mathbb{R}^k} \|\mathbf{v} - Dh(\mathbf{u})^*[\alpha]\|^2 = (Dh(\mathbf{u})^*)^\dagger[\mathbf{v}],$$

where $(Dh(\mathbf{u})^*)^\dagger$ denotes the Moore-Penrose pseudo-inverse of $Dh(\mathbf{u})^*$ [BH12, Section 6].

We can conclude this section by saying that the orthogonal projection of $\mathbf{v} \in \mathbb{R}[x]_d^2$ onto the tangent space to $\mathcal{M}$ at a point $\mathbf{u}$ is obtained by subtracting its projection onto the normal space to $\mathcal{M}$ at $\mathbf{u}$:

$$\text{Proj}_\mathbf{u}(\mathbf{v}) = \mathbf{v} - Dh(\mathbf{u})^* \left[(Dh(\mathbf{u})^*)^\dagger[\mathbf{v}] \right].$$

³By definition of local defining function, the gradients of the h_i at $\mathbf{u}$ are linearly independent as $Dh(\mathbf{u})$ is full rank.

⁴This is known as the adjoint operator of $Dh(\mathbf{u})$ [Bou23, Def 3.4]. Indeed, let $\mathbf{v} \in \mathbb{R}[x]_d^2$ and $\alpha \in \mathbb{R}^k$. By denoting by $\langle \cdot, \cdot \rangle_{\mathbb{R}^k}$ (resp. $\langle \cdot, \cdot \rangle_{\mathbb{R}[x]_d^2}$) the inner product in $\mathbb{R}^k$ (resp. $\mathbb{R}[x]_d^2$),

$$\begin{aligned} \langle Dh(\mathbf{u})[\mathbf{v}], \alpha \rangle_{\mathbb{R}^k} &= \sum_{i=1}^k \alpha_i \langle \text{grad } h_i(\mathbf{u}), \mathbf{v} \rangle_{\mathbb{R}[x]_d^2} \\ &= \left\langle \sum_{i=1}^k \alpha_i \text{grad } h_i(\mathbf{u}), \mathbf{v} \right\rangle_{\mathbb{R}[x]_d^2} \\ &= \langle Dh(\mathbf{u})^*[\alpha], \mathbf{v} \rangle_{\mathbb{R}[x]_d^2}. \end{aligned}$$

6.2 Retraction

The second important concept in the Riemannian Gradient Descent (Algorithm 1) is that of *retraction*. A retraction is a smooth map that, given a point on the tangent space to $\mathcal{M}$ at some point $x \in \mathcal{M}$, returns a point on the manifold $\mathcal{M}$, while satisfying some conditions. More than being a smooth map from $\mathcal{T}_x\mathcal{M}$ to $\mathcal{M}$, it must also vary smoothly with x . Hence, a retraction [Adl+02, Section 3] is defined on the *tangent bundle* to $\mathcal{M}$, which is the disjoint union of the tangent spaces of $\mathcal{M}$ [Bou23, Def 3.42]:

$$\mathcal{TM} := \{(x, v) \mid x \in \mathcal{M} \text{ and } v \in \mathcal{T}_x\mathcal{M}\}.$$

Definition 6.4 – Retraction

A *retraction* on a manifold $\mathcal{M}$ is a smooth map

$$R : \mathcal{V} \rightarrow \mathcal{M} : (x, v) \mapsto R(x, v) = R_x(v),$$

where $\mathcal{V} \subseteq \mathcal{TM}$ is an open subset of $\mathcal{TM}$ and R_x is the restriction of R to $\mathcal{T}_x\mathcal{M} \cap \mathcal{V}$, such that

1. The zero tangent vector at each point of $\mathcal{M}$ belongs to the domain of R :

$$(x, 0_x) \in \mathcal{V} \quad \forall x \in \mathcal{M},$$

where $0_x \in \mathcal{T}_x\mathcal{M}$ denotes the zero tangent vector at $x \in \mathcal{M}$;

2. For all $x \in \mathcal{M}$, $R_x(0_x) = R(x, 0_x) = x$;
3. For all $x \in \mathcal{M}$, $DR_x(0_x) : \mathcal{T}_x\mathcal{M} \rightarrow \mathcal{T}_x\mathcal{M}$ is the identity map.

Therefore, a retraction allows to move away from a point on the manifold along a tangent direction, while remaining on the manifold. Note that, unlike the Riemannian gradient⁵, there may exist several retractions. There is no general theory about what constitutes a good retraction. Sometimes, it may depend on the problem. However, it is preferable that it be cheap to compute, although this often requires a closed form formula, which is not always easy to derive.

In our case, we have not been able to find a closed form retraction for the embedded submanifolds of $\mathbb{R}[x]_d^2$ describing the feasible set of (1.2). We must therefore, for the time being, rely on more general retractions. For example, Absil and Malick showed that the operation consisting of taking the orthogonal projection onto the manifold is a retraction [AM12, Prop 3.2], as well as the orthographic retraction, consisting of mapping a point to the closest point belonging to the intersection between the manifold and the normal space [AM12, Ex. 4.3].

When closed form formulas are not available, we have to rely on optimization algorithms to compute those retractions. This is problematic since it introduces numerical errors (as the optimization process is stopped when some tolerance is reached), while all iterates must lie on the manifold. However, since computers already work with finite representation of numbers, there will always be precision issues, even when we dispose of a closed form

⁵The orthogonal projection onto the tangent space is uniquely defined as it is a convex and closed space.

formula. Hence, we do not expect this to be a big issue⁶.

6.2.1 Newton-like retraction

Another retraction is obtained by applying the Newton-Raphson algorithm to find a zero of a local defining function. The Newton-Raphson algorithm is a well known method for solving nonlinear systems of equations. Given a function $f : \mathbb{R}^m \rightarrow \mathbb{R}^m$ and an initial iterate $x_0 \in \mathbb{R}^m$, it defines a sequence of points $(x_k)_{k \in \mathbb{N}}$ given by

$$x_{k+1} = N_f(x_k) := x_k - [Jf(x_k)]^{-1}f(x_k), \quad (6.1)$$

where $Jf : \mathbb{R}^m \rightarrow \mathbb{R}^{m \times m}$ is the Jacobian⁷ of the function f . Assuming that f is two-time continuously differentiable in a neighbourhood of a root $\bar{x}$ of f , that its Jacobian is nonsingular at $\bar{x}$, and that x_0 is sufficiently close to $\bar{x}$, then the iterates defined by (6.1) converge quadratically to $\bar{x}$ [SM03, Theorem 4.4].

When the input and output spaces of f are not of the same dimension, or when $Jf(\xi)$ is singular at some iterate ξ , the formula has been extended in [Ben66] by using the Moore-Penrose pseudo-inverse [BH12] of the Jacobian, keeping a quadratic convergence rate:

$$x_{k+1} = N_f(x_k) := x_k - [Jf(x_k)]^\dagger f(x_k), \quad k \in \mathbb{N}. \quad (6.2)$$

Consider a smooth manifold $\mathcal{M}$ and a point $u \in \mathcal{M}$ such that there exists a local defining function $h : \mathcal{U} \rightarrow \mathbb{R}^m$ at u for some neighbourhood $\mathcal{U}$ of u . Inspired by [Adl+02], we show that, for points sufficiently close to the manifold, the operation consisting of taking the limit of the Newton-Raphson method's iterates to find a zero of h is a valid retraction.

Theorem 6.5 – Newton-like retraction

Let $\mathcal{M}$ be an embedded submanifold of a Euclidean space $\mathcal{E}$, $x \in \mathcal{M}$, and $h : \mathcal{U} \rightarrow \mathbb{R}^m$ a local defining function at x , with $\mathcal{U} \subseteq \mathcal{E}$ a neighbourhood of x in $\mathcal{E}$.

There exists a neighbourhood $\mathcal{V} \subseteq \mathcal{T}\mathcal{M}$ of $(x, 0_x) \in \mathcal{T}\mathcal{M}$, where 0_x denotes the null vector in $\mathcal{T}_x\mathcal{M}$, such that, for every vector $(\tilde{x}, v) \in \mathcal{V}$, the limit of the sequence of the Newton-Raphson method's iterates applied to $q_0 := \tilde{x} + v$ gives a valid retraction on $\mathcal{M}$:

$$R : \mathcal{V} \rightarrow \mathcal{M} : (\tilde{x}, v) \mapsto \lim_{k \rightarrow \infty} q_k := N_h^k(\tilde{x} + v),$$

where q_k is given by (6.2), and N_h^k denotes k successive applications of N_h . Specifically, we show:

- (a) The restriction R_x of R to $\mathcal{T}_x\mathcal{M}$ applied to the zero vector gives x : $R_x(0_x) = x$;
- (b) The differential of one iteration of N_h at x in a direction $v \in \mathcal{E}$ is the projection of v onto $\mathcal{T}_x\mathcal{M}$: $DN_h(x)[v] = \text{Proj}_{\mathcal{T}_x\mathcal{M}}(v)$. In particular, if $v \in \mathcal{T}_x\mathcal{M}$, this corresponds to the identity map;
- (c) There exists a neighbourhood $\mathcal{W} \subseteq \mathcal{T}_x\mathcal{M}$ of $0_x \in \mathcal{T}_x\mathcal{M}$ such that the differential

⁶Numerical experiments we have conducted worked well, without showing numerical issues due to inaccuracy of the retraction computation.

⁷The *Jacobian* of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ at a point $x \in \mathbb{R}^n$ is defined as the matrix $Jf(x) \in \mathbb{R}^{m \times n}$ such that $Df(x)[v] = Jf(x) \cdot v$ for all $v \in \mathbb{R}^n$.

of the retraction R_x at $0_x \in \mathcal{T}_x\mathcal{M}$ is the identity map: $DR_x(0_x)[w] = w$ for all $w \in \mathcal{W}$.

The smoothness of R is proved in [Bey93].

Proof. (a) If $q_0 = x$, iterates stay constant as it is a zero of h by definition. Therefore,

$$R_x(0) = \lim_{k \rightarrow \infty} q_k = q_0 = x.$$

(b) The differential at x of the Newton operator N_h in a direction $v \in \mathcal{E}$ is given by

$$\begin{aligned} DN_h(x)[v] &= v - D\left([Jh(x)]^\dagger \cdot h(x)\right)[v] \\ &= v - D\left([Jh(x)]^\dagger\right)[v] \cdot h(x) - [Jh(x)]^\dagger \cdot Dh(x)[v] \\ &\quad \downarrow \text{Using the facts that } h(x) = 0 \text{ and } Dh(x)[v] = Jh(x) \cdot v \\ &= v - [Jh(x)]^\dagger \cdot Jh(x) \cdot v \\ &\quad \downarrow \text{The differential of a local defining function being surjective (Definition 4.3)} \\ &= v - [Jh(x)]^T \cdot \left(Jh(x) \cdot [Jh(x)]^T\right)^{-1} \cdot Jh(x) \cdot v \\ &= v - [Jh(x)]^T \cdot \left([Jh(x)]^T\right)^\dagger \cdot v \\ &\quad \downarrow \text{By definition of the normal space to } \mathcal{M} \text{ at } x \text{ (see Section 6.1.2)} \\ &= v - \text{Proj}_{\mathcal{N}_x\mathcal{M}}(v) = \text{Proj}_{\mathcal{T}_x\mathcal{M}}(v). \end{aligned}$$

(c) By recursion, we are now able to show that for all $k \in \mathbb{N}_0$,

$$DN_h^k(x)[v] = \text{Proj}_{\mathcal{T}_x\mathcal{M}}(v).$$

Assume it is valid for some $k \in \mathbb{N}_0$ (the case $k = 1$ has just been proved). Then,

$$\begin{aligned} DN_h^k(x)[v] &= DN_h(N_h^{k-1}(x))[DN_h^{k-1}(x)[v]] \\ &= DN_h(x)[v] = v. \end{aligned}$$

This allows us to compute the differential of R_x at $0 \in \mathcal{T}_x\mathcal{M}$ in the direction $w \in \mathcal{W}$, for some neighbourhood $\mathcal{W} \subseteq \mathcal{T}_x\mathcal{M}$ of 0:

$$\begin{aligned} DR_x(0)[w] &= D\left(\lim_{k \rightarrow \infty} N_h^k\right)(x)[w] \\ &= \lim_{k \rightarrow \infty} DN_h^k(x)[w] = \lim_{k \rightarrow \infty} w = w. \end{aligned}$$

In the second equality, we swapped the limit and differential operators, which is valid if there exists a neighbourhood of x such that the sequence of derivatives converges uniformly on the neighbourhood and there exists a point in the neighbourhood for which the sequence converges [BS11]. The uniform convergence of the derivatives is shown in [Bey93]. For the second condition, we have the convergence of the sequence at x , for which the sequence converges to x . $\square$

Conclusion

Nonnegativity of polynomials is and remains a major concern of mathematics to this day. While we do not expect to characterize exactly the set of nonnegative polynomials in the general case (unless $P = NP$), we often rely on the relaxation of sum-of-squares polynomials.

This Master's thesis aimed to make the first crack in the wall of ice separating us from a fully understanding of the structure of linear sum-of-squares problems of the kind (1.2). While feasibility problems have already been studied in [LYP23], the novel part of our work relies in the introduction of a linear objective in the sum of squares.

Problem (1.2) is a nonlinear relaxation of a more general kind of problems, described by (1.1), in which we search polynomials that can be expressed as the sum of two squared polynomials. For simplicity, we considered only univariate polynomials. In this case, being nonnegative is equivalent to being the sum of two squares (Lemma 2.2), and the reformulation is exact. The interest of the SOS reformulation (1.2) lies in the small number of variables it contains, in comparison with the original problem (1.1), which can be expressed as an SDP but whose number of variables grows quadratically with the size of the problem.

Yet, it introduces nonconvexity and possibly non smoothness in the feasible set. Throughout the chapters, we analyzed increasingly complex feasible sets to understand their structure. In Chapter 3, we showed that the set of pairs of coprime polynomials $\mathbb{R}[x]_d^{2,c}$ of a given degree is a smooth embedded submanifold of the vector space $\mathbb{R}[x]_d^2$. Further, in Chapter 4, we gave an example of embedded submanifold, corresponding to a feasible set defined by a bijective linear constraint operator $\mathcal{L}_2$. Next, we showed that, at a feasible point corresponding to a pair of coprime polynomials, the feasible set is locally smooth. Finally, in Chapter 5, we developed an implicit expression of the tangent cone to the feasible set at each of its points. At a feasible point $\mathbf{u}$ such that the kernel of the constraint operator $\mathcal{L}_2$ is contained in the image of the operator $\mathcal{A}_u$, we showed that the set is locally smooth. However, around a feasible point $\mathbf{u}$ such that the image of $\mathcal{A}_u$ is strictly contained in the kernel of $\mathcal{L}_2$, the set is not smooth. In fact, if such a point exists, the feasible set seems to be a stratified set.

When the feasible set is a smooth manifold, we may use a Riemannian Gradient Descent method to find an optimum of our main problem (1.2), which is described in Chapter 6. For this, we detailed the computation of the Riemannian gradient at a point, and a possible choice of retraction. For this retraction, we prove that applying a Newton-Raphson algorithm to find a zero of a local defining function for the manifold at the considered point, is a valid retraction. This result may be of independent interest, as it is also valid on other smooth manifolds than the one of pairs of polynomials. While this may introduce errors due to the numerical imprecisions of the resulting point (obtained after

an optimization process stopped when some tolerance is reached), we do not expect this to be a significant concern, as computers work with finite precision anyway.

While the results obtained are already promising, there remain some unexplained cases, starting with the one of a feasible set containing a point $\mathbf{u}$ at which the intersection between the image of $\mathcal{A}_{\mathbf{u}}$ and the kernel of $\mathcal{L}_2$ is strictly smaller than the two sets (i.e., neither of them is contained in the other). Once full understanding of the univariate case has been reached, we will be ready to analyze multivariate polynomials.

Behind this work, there are also two hidden goals. The first one, more practical, is the development of an efficient algorithm for solving such problems. For this, we already gave a starting point by describing the Riemannian Gradient Descent in Chapter 6. Yet, if we could obtain a retraction with a closed form formula, it would prevent us from applying a nested optimization algorithm inside the Riemannian Gradient Descent, improving (theoretically) the efficiency. The second goal, more theoretical, is to understand the global landscape of the problem, and whether it has benign properties we could exploit. One such “popular” property is benign nonconvexity, i.e., nonconvex problems in which all the local minima are in fact global. Combining ideas of Legat et al. [LYP23] and of Boumal et al. [BVB18], we can expect to describe the set of linear sum-of-squares problems for which we have benign nonconvexity. In the first paper, the authors show that the problem of finding a feasible point to our main problem (1.2) can be expressed as a minimization of a quadratic penalty, which exhibits this benign nonconvexity. In the second one, the authors show that, if the feasible set of the Burer-Monteiro reformulation of an SDP problem has a smooth structure, we have once again benign nonconvexity for almost all linear objectives. As our work combines polynomials and smooth manifolds, we expect to be able to show the benign nonconvexity of problem (1.2) by merging the ideas of both papers, with particular focus on the cases where the feasible set is not a smooth manifold.

Bibliography

- [Adl+02] Roy L. Adler et al. “Newton’s Method on Riemannian Manifolds and a Geometric Model for the Human Spine”. In: *IMA Journal of Numerical Analysis* 22.3 (July 1, 2002), pp. 359–390.
- [AM12] Pierre-Antoine Absil and Jérôme Malick. “Projection-like Retractions on Matrix Manifolds”. In: *SIAM Journal on Optimization* 22.1 (Jan. 2012), pp. 135–158.
- [AM16] Amir Ali Ahmadi and Anirudha Majumdar. “Some Applications of Polynomial Optimization in Operations Research and Real-Time Decision Making”. In: *Optimization Letters* 10.4 (Apr. 1, 2016), pp. 709–729.
- [AMS08] Pierre-Antoine Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton, NJ: Princeton University Press, 2008.
- [Ani16] Majumdar Anirudha. “Funnel Libraries for Real-Time Robust Feedback Motion Planning”. Department of Electrical Engineering and Computer Science: Massachusetts Institute of Technology, 2016.
- [Ano] Anonymous. *SIAM Activity Group on Optimization Test of Time Award*. Society for Industrial and Applied Mathematics. URL: <https://www.siam.org/programs-initiatives/prizes-awards/activity-group-prizes/siam-activity-group-on-optimization-test-of-time-award/> (visited on 05/09/2026).
- [AP15] James Anderson and Antonis Papachristodoulou. “Advances in Computational Lyapunov Analysis Using Sum-of-Squares Programming.” In: *Discrete and Continuous Dynamical Systems*. B 20 (Aug. 2015), pp. 2361–2381.
- [Bai89] Anatole Bailly. *Apocalupsis*. In: *Dictionnaire Grec - Français*. 41st ed. Paris: Hachette, 1989, p. 226.
- [Bar95] A. I. Barvinok. “Problems of Distance Geometry and Convex Properties of Quadratic Maps”. In: *Discrete & Computational Geometry* 13.2 (Mar. 1, 1995), pp. 189–202.
- [Ben66] Adi Ben-Israel. “A Newton-Raphson Method for the Solution of Systems of Equations”. In: *Journal of Mathematical Analysis and Applications* 15.2 (Aug. 1, 1966), pp. 243–252.
- [Bey93] W. -J. Beyn. “On Smoothness and Invariance Properties of the Gauss-Newton Method”. In: *Numerical Functional Analysis and Optimization* 14.5–6 (Jan. 1, 1993), pp. 503–514.

- [BH12] J. C. A. Barata and M. S. Hussein. “The Moore-Penrose Pseudoinverse. A Tutorial Review of the Theory”. In: *Brazilian Journal of Physics* 42.1–2 (Apr. 2012), pp. 146–165.
- [Ble+24] Grigoriy Blekherman et al. *Spurious Local Minima in Nonconvex Sum-of-Squares Optimization*. Nov. 4, 2024. URL: <http://arxiv.org/abs/2411.02208> (visited on 01/14/2026). Pre-published.
- [BM03] Samuel Burer and Renato D.C. Monteiro. “A Nonlinear Programming Algorithm for Solving Semidefinite Programs via Low-Rank Factorization”. In: *Mathematical Programming* 95.2 (Feb. 1, 2003), pp. 329–357.
- [BM05] Samuel Burer and Renato D.C. Monteiro. “Local Minima and Convergence in Low-Rank Semidefinite Programming”. In: *Mathematical Programming* 103.3 (July 1, 2005), pp. 427–444.
- [Bou23] Nicolas Boumal. *An Introduction to Optimization on Smooth Manifolds*. 1st ed. Cambridge University Press, Mar. 16, 2023.
- [Bou26] Nicolas Boumal. *A Different Way to Define Embedded Submanifolds: Part I*. the OPTIM@EPFL blog. Feb. 10, 2026. URL: <https://www.racetothetobottom.xyz/posts/smooth-retracts-are-manifolds-part1/> (visited on 04/28/2026).
- [BPT12] Grigoriy Blekherman, Pablo A. Parrilo, and Rekha R. Thomas. *Semidefinite Optimization and Convex Algebraic Geometry*. SIAM, 2012.
- [BS11] Robert G. Bartle and Donald R. Sherbert. *Introduction to Real Analysis*. 4th ed. University of Illinois: John Wiley & Sons, Inc., Jan. 18, 2011.
- [BV04] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [BVB18] Nicolas Boumal, Vladislav Voroninski, and Afonso S. Bandeira. *The Non-Convex Burer-Monteiro Approach Works on Smooth Semidefinite Programs*. Apr. 10, 2018. URL: <http://arxiv.org/abs/1606.04970> (visited on 01/14/2026). Pre-published.
- [CLO05] David A. Cox, John Little, and Donal O’Shea. *Using Algebraic Geometry*. 2nd ed. Vol. 185. Graduate Texts in Mathematics. New York: Springer, 2005.
- [CLO25] David A. Cox, John Little, and Donal O’Shea. *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*. Undergraduate Texts in Mathematics. Cham: Springer Nature Switzerland, 2025.
- [CLR94] Man-Duen Choi, Tsit Lam, and B. Reznick. “Sums of Squares of Real Polynomials”. In: *Proceedings of Symposia in Pure Mathematics* 58 (Jan. 1, 1994), pp. 103–126.
- [Emb16] Mark Embree. *Lecture Notes on Numerical Analysis*. Virginia Tech - MATH/CS 5466, Spr. 2016.
- [Far25] Gabriele Farina. *Polynomial Optimization*. Lecture notes for MIT 6.7220. MIT, Spr. 2025.
- [Hil00] David Hilbert. “Mathematische Probleme”. In: *Nachrichten von der Königlichen Gesellschaft der Wissenschaften zu Göttingen* (1900), pp. 253–297.

- [HL05] Didier Henrion and Jean-Bernard Lasserre. “Detecting Global Optimality and Extracting Solutions in GloptiPoly”. In: *Positive Polynomials in Control*. Berlin, Heidelberg: Springer, 2005, pp. 293–310.
- [Kha22] Hassan K. Khalil. “Lyapunov Stability”. In: *Nonlinear Systems*. 3rd ed. Michigan State University: Prentice Hall, 2022, pp. 111–194.
- [Las01] Jean-Bernard Lasserre. “Global Optimization with Polynomials and the Problem of Moments”. In: *SIAM Journal on Optimization* 11 (2001), pp. 796–817.
- [Las10] Jean-Bernard Lasserre. *Moments, Positive Polynomials and Their Applications*. Vol. 1. Imperial College Press Optimization Zeries. London: Imperial College Press, 2010.
- [Lee12] John M. Lee. *Introduction to Smooth Manifolds*. Vol. 218. Graduate Texts in Mathematics. New York, NY: Springer, 2012.
- [LKB23] Eitan Levin, Joe Kileel, and Nicolas Boumal. “Finding Stationary Points on Bounded-Rank Matrices: A Geometric Hurdle and a Smooth Remedy”. In: *Mathematical Programming* 199.1 (May 1, 2023), pp. 831–864.
- [LYP23] Benoît Legat, Chenyang Yuan, and Pablo A. Parrilo. “Low-Rank Univariate Sum of Squares Has No Spurious Local Minima”. In: *SIAM Journal on Optimization* 33.3 (Sept. 30, 2023), pp. 2041–2061.
- [Nes00] Yurii Nesterov. “Squared Functional Systems and Optimization Problems”. In: *High Performance Optimization*. Ed. by Hans Frenk et al. Boston, MA: Springer US, 2000, pp. 405–440.
- [Nes18] Yurii Nesterov. *Lectures on Convex Optimization*. Vol. 137. Springer Optimization and Its Applications. Cham: Springer International Publishing, 2018.
- [NW06] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. 2nd ed. New York: Springer, Dec. 11, 2006.
- [OGA23] Guillaume Olikier, Kyle A. Gallivan, and P.-A. Absil. *First-Order Optimization on Stratified Sets*. Mar. 28, 2023. URL: <http://arxiv.org/abs/2303.16040> (visited on 03/11/2026). Pre-published.
- [Oli+25] Guillaume Olikier et al. *The Tangent Cone to the Real Determinantal Variety: Various Expressions and a Proof*. June 7, 2025. URL: <http://arxiv.org/abs/2504.11382> (visited on 01/12/2026). Pre-published.
- [Par00] Pablo A. Parrilo. “Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization”. PhD thesis. California Institute of Technology, 2000.
- [Pat98] Gábor Pataki. “On the Rank of Extreme Matrices in Semidefinite Programs and the Multiplicity of Optimal Eigenvalues”. In: *Mathematics of Operations Research* 23.2 (May 1998), pp. 339–358.
- [Rez00] Bruce Reznick. “Some Concrete Aspects of Hilbert’s 17th Problem”. In: *Real algebraic geometry and ordered structures* 253 (2000).
- [RW98] “Variational Geometry”. In: *Variational Analysis*. Ed. by R. Tyrrell Rockafellar and Roger J. B. Wets. Berlin, Heidelberg: Springer, 1998, pp. 196–237.

- [SM03] Endre Süli and David F. Mayers. “Simultaneous Nonlinear Equations”. In: *An Introduction to Numerical Analysis*. Cambridge: Cambridge University Press, 2003, pp. 104–132.
- [Str16] Gilbert Strang. *Introduction to Linear Algebra*. 5th ed. Wellesley: Cambridge Press, 2016.
- [Tes19] Gerald Teschl. “Metric and Topological Spaces”. In: *Topics in Real and Functional Analysis*. Providence, Rhode Island: American Mathematical Society, Jan. 14, 2019, pp. 529–556.
- [Wee06] Tan Weehong. “Nonlinear Control Analysis and Synthesis Using Sum-of-Squares Programmin”. Berkeley: University of California, 2006.
- [WM22] John Wright and Yi Ma. “Augmented Lagrange Multipliers”. In: *High-Dimensional Data Analysis with Low-Dimensional Models: Principles, Computation, and Applications*. Cambridge University Press, 2022, pp. 331–338.
- [Woo16] Henry Woody. *Polynomial Resultants*. Department of Mathematics and Computer Science University of Puget Sound, 2016.
- [Yan+23] Lujie Yang et al. “Approximate Optimal Controller Synthesis for Cart-Poles and Quadrotors via Sums-of-Squares”. In: (2023).
- [Yan23] Heng Yang. *Optimal Control and Estimation*. ES/AM 158. Harvard, Nov. 14, 2023.
- [Yan24] Heng Yang. *Semidefinite Optimization and Relaxation*. ENG-SCI 257. Harvard, Apr. 15, 2024.
- [Yua22] Chenyang Yuan. “Polynomial Structure in Semidefinite Relaxations and Non-Convex Formulations”. Thesis. Massachusetts Institute of Technology, Sept. 2022.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
École polytechnique de Louvain

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