

OPTIMIZATION MODELS FOR SHAPE-CONSTRAINED FUNCTION
ESTIMATION PROBLEMS INVOLVING NONNEGATIVE POLYNOMIALS AND
THEIR RESTRICTIONS

by
DÁVID PAPP

A Dissertation submitted to the
Graduate School-New Brunswick
Rutgers, The State University of New Jersey
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy
Graduate Program in Operations Research
written under the direction of
Dr. Farid Alizadeh
and approved by

New Brunswick, New Jersey

May 2011

ABSTRACT OF THE DISSERTATION

Optimization models for shape-constrained function estimation problems involving nonnegative polynomials and their restrictions

by DÁVID PAPP

Dissertation Director:

Dr. Farid Alizadeh

In this thesis function estimation problems are considered that involve constraints on the shape of the estimator or some other underlying function. The function to be estimated is assumed to be continuous on an interval, and is approximated by a polynomial spline of given degree.

First the estimation of univariate functions is considered under constraints that can be represented by the nonnegativity of affine functionals of the estimated function. These include bound constraints, monotonicity, and convexity constraints. A general framework is presented in which arbitrary combination of such constraints can be modeled and handled in a both theoretically and practically efficient manner, using semidefinite programming. The approach is illustrated by a variety of applications in statistical estimation. Numerical results are compared to those obtained using methods previously proposed in the statistical literature.

Next, multivariate problems are considered. Shape constraints that are tractable in the univariate case are intractable in the multivariate setting. Hence, instead of following the univariate approach, polyhedral and spectrahedral (semidefinite) approximations are proposed, which are obtained by replacing nonnegative splines by piecewise weighted-sum-of-squares polynomial splines. A decomposition method for solving very large scale multivariate problems is also investigated.

In order to include more constraints in the same unified framework, the notion of weighted-sum-of-squares functions is generalized to weighted-sum-of-squares cones in finite-dimensional real linear spaces endowed with a bilinear multiplication. It is shown that such cones are semidefinite representable, and the representation is explicitly constructed.

Finally, seeking a method to circumvent the use of semidefinite programming in the solution of optimization problems over nonnegative and weighted-sum-of-squares polynomials, we study the conditions of optimality, as well as barrier functions in such optimization problems. For univariate polynomials, a method based on fast Fourier transform is provided that accelerates the evaluation of the logarithmic barrier function (and its first two derivatives) of the moment cone. This reduces

the running time of some interior point methods for semidefinite programming from $\mathcal{O}(n^6)$ to $\mathcal{O}(n^3)$ when solving problems involving polynomial nonnegativity constraints.

Contents

Abstract	ii
1 Introduction	1
1.1 Outline of the thesis	5
2 Nonnegative and sum-of-squares polynomials	7
2.1 Univariate nonnegative polynomials	9
2.2 Multivariate nonnegative polynomials	13
2.2.1 Global nonnegativity	13
2.2.2 Nonnegative polynomials on polyhedra	15
2.3 Dual cones of nonnegative polynomials: moment cones	18
3 Univariate shape constrained estimation	21
3.1 Existing results in univariate shape-constrained estimation	22
3.2 Sieves and cones of nonnegative functions	26
3.2.1 Estimation in functional Hilbert spaces	26
3.2.2 Sieves and finite-dimensional approximations	28
3.3 Representations of nonnegative polynomials and splines	32
3.3.1 Representations involving semidefinite and second-order cone constraints	32
3.3.2 Sieves of polyhedral spline cones	34
3.4 Shape-constrained statistical estimation models	38
3.4.1 Non-parametric regression of a nonnegative function	39
3.4.2 Isotonic, convex, and concave non-parametric regression	39
3.4.3 Unconstrained density estimation	40
3.4.4 Unimodal density estimation	40

3.4.5	Log-concave density estimation	41
3.4.6	Arrival rate estimation	44
3.4.7	Knot point selection	45
3.5	Numerical examples	46
3.5.1	Isotonic and convex/concave regression	46
3.5.2	Density estimation	51
3.5.3	Arrival rate estimation	53
3.6	Summary and conclusion	54
4	Multivariate shape constrained estimation	56
4.1	Multivariate nonnegative and sum-of-squares splines	57
4.1.1	Numerical illustration – posterior probability estimation	61
4.2	A decomposition method for multivariate splines	63
4.2.1	Augmented Lagrangian decomposition for sparse problems	64
4.2.2	Numerical illustration – NJ Turnpike accidents	67
4.3	Generalized sum-of-squares cones	70
4.3.1	Semidefinite characterization of sums of squares	72
4.3.2	Application to vector-valued functions	78
4.3.3	Further applications	82
4.4	Conclusion	84
5	Direct optimization methods for nonnegative polynomials	85
5.1	Optimality conditions for nonnegative polynomials	87
5.1.1	Preliminary observations	89
5.1.2	Nonnegative polynomials over a finite interval	91
5.1.3	Müntz polynomials	100
5.1.4	Cones linearly isomorphic to nonnegative polynomials over an interval	102
5.2	A barrier method for moment cone constraints	106
5.2.1	Fast computation of the gradient and the Hessian	107
5.2.2	Alternative bases	109
5.3	Conclusion	110
6	Conclusion	112

A Background material	114
A.1 Conic optimization	114
A.2 Positive semidefinite matrices	116
A.3 Semidefinite and second-order cone programming	117
Bibliography	120
Vita	131

Chapter 1

Introduction

In this thesis we consider function estimation problems involving constraints on the shape of the estimator or some other underlying function. Such estimation problems abound in statistics and engineering, and there are numerous examples in other disciplines as well. Here we briefly recall a few to give a flavor of these applications; we will return to some of them in the subsequent chapters.

1. **Isotonic regression.** Regression problems often arise in situations where some prior knowledge (laws of nature or common sense) tells us that the estimated function must be monotone increasing or decreasing in some of the variables. Calculating the least-squares regression function under this constraint is the simplest form of the univariate *isotonic regression* problem [RWD88]. Multivariate extensions, as well as more abstract forms of this problem, such as [Bru55], also exist; in these problems monotonicity is understood with respect to some partial order.
2. **Convex/Concave regression.** Similarly to isotonic regression, some problems may require that the estimator is convex or concave over its domain [Hil54]. Univariate applications include the estimation of the Laffer curve [Wan78] and utility functions in econometrics. The latter is also an example of a situation when multiple shape constraints may be present: the utility function of a rational, risk-averse decision maker is increasing and concave.
3. **Density estimation.** One of the fundamental problems in statistics is the estimation of probability density functions (or PDFs) and cumulative distribution functions (or CDFs) of random variables, given a finite number of samples (realizations) of these variables [DG85, EL01]. These problems come in two flavors: in the parametric setting the unknown PDF or CDF is

assumed to belong to a finite-dimensional space (or at least to a family of functions given up to a finite number of parameters) that contains only PDFs and CDFs. Usually one considers parametric families which contain only functions of the right shape. In the non-parametric setting, however, no such assumptions are made, and constraints such as nonnegativity, being less than or equal to one, and monotonicity must be explicitly added to the problem. Further constraints, such as the unimodality or log-concavity of the PDF can also be considered (even though unimodality is not a convex constraint, which makes optimization with unimodality constraints particularly difficult).

4. **Arrival rate estimation.** The arrival rate of a non-homogeneous Poisson process must be nonnegative by definition. The paper [AENR08] is the first in which a similar non-parametric method to the one in this thesis is used in the univariate case. See also [PA08] and the thesis [Rud09].
5. **Design of curves.** Various engineering problems involve the design of curves with shape constraints [BF84, GSY94]. In particular, contours of objects and curves describing controlled moving objects (such as a robot arm) may be required to lie entirely within a region. The region may be, for example, a polyhedral set, or a basic semialgebraic set defined by polynomial inequalities. The curve along which objects are moved may have a restricted curvature.
6. **Estimation of posterior probabilities.** Suppose a sample is known to belong to one of N populations, and we wish to estimate the probability that it was sampled from population $i \in \{1, \dots, N\}$, given some information on the distribution of the individual populations. When $N = 2$, and one of the two classes is “good” while the other is “bad”, this can be viewed as the assignment of a risk score to the sample. Such models are abound in a number of areas ranging from medicine to homeland security.
7. **Stochastic dominance.** One way to express a decision maker’s preferences among different gambles, or more specifically, to introduce risk-aversion in stochastic optimization models is using *stochastic dominance* constraints [HR69]. We say that (real-valued) random variable A has *first-order dominance* over B if their CDFs F_A and F_B satisfy $F_A(t) \leq F_B(t)$ for every t . In models B is often a parameter, a “benchmark”. First-order dominance can be viewed as a variant of the requirement that the graph of the estimated function lies in a given region. A weaker form of dominance, *second-order dominance*, is expressible as $\int_{-\infty}^s F_A(t)dt \leq \int_{-\infty}^s F_B(t)dt$ for every s . Optimization problems involving stochastic dominance constraints were introduced in

[DR03], and were developed in a series of papers by the same authors.

8. **Covariance matrix estimation.** Covariance matrix estimation in time series can be used to capture the evolution of interdependencies between multiple time series. The object of interest is then a univariate matrix-valued function that is required to be positive semidefinite throughout the considered time interval. Depending on the estimation method this constraint need not be satisfied automatically at every point by the estimator, in which case it needs to be included as a hard constraint.

Perhaps the simplest of shape constraints in the above examples is nonnegativity, or (only slightly more generally), the requirement that the graph of the estimated function is on one side of a given line. Many other shape constraints, under some differentiability assumptions, also reduce to the same constraint. For example, monotonicity, convexity, and concavity of (twice) differentiable functions can be expressed as a constraint on the sign of the first or second derivative of the function. Hence, a good starting point for our investigations is the following general optimization problem:

$$\begin{aligned} & \text{minimize} && c(\mathbf{x}) \\ & \text{subject to} && \mathcal{A}_i(f_{\mathbf{x}})(\mathbf{t}) - b_i \geq 0 \quad \forall \mathbf{t} \in \Delta, \quad i = 1, \dots, k, \\ & && \mathbf{x} \in \mathcal{X}, \end{aligned} \tag{1.1}$$

where \mathcal{X} is an arbitrary (albeit usually convex, close, and finite-dimensional) set; $f_{\mathbf{x}}: \Delta \rightarrow \mathbb{R}$ is a real-valued function with domain $\Delta \subseteq \mathbb{R}^m$ for every \mathbf{x} ; c is a (usually convex) cost function, each \mathcal{A}_i is a linear operator, and $b_i \in \mathbb{R}$.

The minimization in (1.1) is with respect to the vector \mathbf{x} . With a slight abuse of notation we can consider the estimated function $f_{\mathbf{x}}$ to be the variable. However, it is often helpful to distinguish between the unknown *function* $f_{\mathbf{x}}$ and its *representation* \mathbf{x} . Most algorithmic results, for example, do not depend on the properties of the set of feasible functions $\{f_{\mathbf{x}} \mid \mathbf{x} \in \mathcal{X}\}$, but rather on those of the set \mathcal{X} .

In most of the applications considered in this thesis, \mathcal{X} is a finite-dimensional set, while Δ is a polyhedral set. (Some of the results generalize to the setting where Δ is a basic semialgebraic set.) In these cases, (1.1) can be cast as a semi-infinite optimization problem (and if c is affine, then as a semi-infinite linear optimization problem) with a continuum of linear constraints; see [GL98, RR98] for an extensive review of semi-infinite linear optimization methods. In search of more efficient methods, in this thesis we do not follow this direction. Instead, we assume that $f_{\mathbf{x}}$ belongs

to a set of functions, such as polynomials of a given degree, polynomial splines of a given order, or trigonometric polynomials of some degree, in which nonnegativity over the entire domain can be effectively characterized or approximated. Our primary models will involve polynomial splines.

Definition 1.1. A function $f: [a, b] \rightarrow \mathbb{R}$ is a *univariate polynomial spline* of degree d and continuity \mathcal{C}^r ($r \leq d - 1$) if there exist *knot points* a_0, \dots, a_m satisfying $a = a_0 < \dots < a_m = b$ such that (1) f is a polynomial of degree d over each interval $[a_i, a_{i+1}]$; (2) f has continuous derivatives up to order r over (a, b) .

In this thesis we will always assume $r = d - 1$, such splines are often referred to as splines with *simple knots*. Schumaker [Sch81] has a more general definition which allows a spline to be differentiable up to different orders at different knot points. Wahba [Wah90] and others consider only natural splines, which are more restricted on the first and last subinterval of the domain. All models in this thesis could be adapted if we were to use any of these definitions. In the multivariate models *tensor product splines* will be used, which are constructed by taking the tensor product of univariate spline spaces.

Not all natural shape constraints can be cast in the form (1.1), and in some situations such a formulation may not be the most useful. Consider, for example, the convexity of multivariate functions. A twice continuously differentiable real-valued function f is convex over $\Delta \subseteq \mathbb{R}^n$ if and only if its Hessian \mathbf{H} is positive semidefinite over Δ . While this condition is equivalent to the nonnegativity of the $2n$ -dimensional function $(\mathbf{x}, \mathbf{y}) \rightarrow \mathbf{y}^T \mathbf{H}(\mathbf{x}) \mathbf{y}$ over $\Delta \times \mathbb{R}^n$, this approach has several drawbacks. A similar constraint arises in the covariance matrix estimation problem. Another example is the curve design problem with curvature constraint. For this problem, no formulation of the form (1.1) is known. We will return to these problems in Chapter 4. At this point it suffices to note that the following generalization of (1.1) can be used to model all of the aforementioned shape constraints:

$$\begin{aligned} & \text{minimize} && c(\mathbf{x}) \\ & \text{subject to} && \mathcal{A}_i(f_{\mathbf{x}})(\mathbf{t}) - b_i \in \mathcal{K} \quad \forall \mathbf{t} \in \Delta, \quad i = 1, \dots, k, \\ & && \mathbf{x} \in \mathcal{X}, \end{aligned} \tag{1.1'}$$

where \mathcal{K} is a closed, convex, and pointed cone. (See Appendix A for definitions.) If $\mathcal{K} = \mathbb{R}$ in (1.1'), we obtain (1.1). The multivariate convexity constraint will be modeled by choosing \mathcal{K} to be the cone of $n \times n$ positive semidefinite real symmetric matrices.

The thesis is self-contained; for completeness, some essential background material on conic and semidefinite programming is included in the Appendix.

On notation. We will use the following notations and conventions: vectors (and matrices) are typeset boldface, their components (rows and columns) are denoted with the corresponding lowercase italic character, and they are indexed starting from 0 rather than 1, as this is more convenient for coefficient vectors of polynomials. For example, the $n + 1$ dimensional row vector \mathbf{p}^T could also be written as (p_0, \dots, p_n) . The inequality $\mathbf{X} \succcurlyeq 0$ denotes that \mathbf{X} is a positive semidefinite real symmetric matrix. All the remaining notation is introduced in the chapters as needed.

1.1 Outline of the thesis

Chapter 2 is primarily a literature review summarizing known characterizations of univariate non-negative polynomials and subsets of multivariate polynomials (“sum-of-squares” and “weighted-sum-of-squares” polynomials) that admit a good characterization suitable for convex optimization models. These characterizations are the most important tools used in the rest of the thesis. We shall add a few improvements and generalizations to the known theorems, as well as some new elementary proofs to some of them. A brief overview of moment cones (the dual cones of nonnegative polynomials) is also included, as it is necessary for Chapter 5.

In Chapter 3 we study univariate function estimation problems with nonnegativity, monotonicity, convexity, and log-concavity constraints. The proposed optimization models involve nonnegative polynomial splines, and can be translated to semidefinite and second-order cone programming problems using the techniques introduced in Chapter 2. We investigate polyhedral approximations to these problems, and give a necessary and sufficient condition for a family of piecewise nonnegative polynomial splines to be dense (in uniform norm) in the cone of nonnegative continuous functions (Theorem 3.4). Specifically, we make a case for using piecewise and coordinate-wise Bernstein polynomial splines. This is followed by an extensive list of optimization models (Section 3.4) for a variety of (statistical) estimation problems, and numerical experiments (Section 3.5), in which the proposed semidefinite programming and linear programming approaches are compared to each other and to kernel methods.

In Chapter 4 we investigate how the ideas of Chapter 3 can be applied in multivariate shape-constrained estimation. Unlike in the univariate setting, nonnegative multivariate polynomials do not

seem to admit a good characterization (their recognition is NP-hard). Nevertheless, using a multivariate generalization of the density theorem of the previous chapter (Theorem 4.1) we provide a method for estimating multivariate nonnegative continuous functions that is both practical and theoretically sound. In Section 4.2 we demonstrate how these techniques can be coupled with decomposition methods to solve the large-scale optimization problems that arise in multivariate shape-constrained spline estimation. Finally, Section 4.3 contains a generalization of Nesterov’s characterization of real-valued sum-of-squares functions (Proposition 2.1) to arbitrary finite-dimensional real linear spaces, where squaring is with respect to an arbitrary bilinear multiplication (Theorem 4.5). A weighted-sum-of-squares version is also provided. These can be applied to a number of shape-constrained optimization problems that cannot be modeled using nonnegativity constraints alone.

Chapter 5 is concerned with approaches that could be used to circumvent the use of semidefinite programming in optimization problems involving nonnegative polynomial constraints. In Section 5.1 we study the conditions of optimality for the cone of nonnegative polynomials with the goal of characterizing complementary pairs of nonnegative polynomials and moment vectors in a way that is suitable for devising homotopy methods for these cones. The main results of this section are negative. In Section 5.2 a barrier method is given that has the same iteration complexity as semidefinite programming for nonnegative polynomial constraints, but whose time complexity is $\mathcal{O}(n^3)$ compared to the $\mathcal{O}(n^6)$ complexity of the semidefinite programming-based method.

In Chapter 6 we summarize our results and discuss some open questions. The Appendix gives some background on semidefinite programming, second-order cone programming, and general conic programming.

All proofs of the Theorems and Lemmas in this thesis are original works of the author. Necessary theorems and lemmas from the literature are reproduced here as Propositions, regardless of their depth, so that they are clearly distinguished from new results. The proofs of previous results are usually omitted, unless their ideas are important for the thesis; the original works in which these results were proven are cited. While many of the new results were obtained during joint work with several co-authors, only those parts of the joint work are reproduced here to which the author of the thesis was the main contributor.

Parts of this thesis have been published in peer-reviewed journals, or have been available as technical reports. Parts of Chapter 3 have been disseminated in [PA08], but in the current, considerably improved, form it is currently under review [PA]. Section 4.3 is currently under review [PCA], while Section 5.1 is part of the accepted paper [RNPA]. The rest of Chapters 4 and 5 first appear here.

Chapter 2

Nonnegative and sum-of-squares polynomials

Optimization problems involving nonnegative and sum-of-squares polynomials will be used as a fundamental modeling tool throughout the thesis. This chapter is a brief summary of the main results about characterizations of nonnegative polynomials, which are necessary to make this thesis self-contained. For a more comprehensive summary, especially regarding the applications of this theory in other fields, such as algebraic geometry and global optimization of polynomials, the reader is advised to consult the excellent monographs [Mar08] or [Lau08]. Unlike in the other chapters, we include some of the prior results with new proofs when these proofs serve as new elementary proofs to some well-known theorems, or when they generalize to some previously unknown results used in subsequent chapters.

We say that a polynomial is *sum-of-squares* if it can be expressed as a sum of finitely many perfect squares. The cone of k -variate sum-of-squares polynomials of total degree $2n$ is denoted by $\Sigma_{2n,k}$.

Let $\mathcal{P}_{2n+1} \in \mathbb{R}^{2n+1}$ denote the cone of those univariate polynomials of degree $2n$ which are nonnegative over the entire real line. Similarly, let \mathcal{P}_{n+1}^Δ denote the cone of those univariate polynomials of degree n which are nonnegative over Δ . Although it is a slight abuse of notation (recall the distinction between \mathbf{x} and $f_{\mathbf{x}}$ in the paragraph following (1.1)), we will also denote the cone of the *coefficient vectors* of nonnegative polynomials (in standard basis, for simplicity) by the same symbol:

$$\mathcal{P}_{n+1}^\Delta = \left\{ (p_0, \dots, p_n)^\top \mid \sum_{i=0}^n p_i t^i \geq 0 \quad \forall t \in \Delta \right\}.$$

The subscript $n + 1$ shows the dimension of the cone. In this thesis the domain Δ is always the entire real line, or an interval. In the multivariate setting $\mathcal{P}_{n,k}^\Delta$ denotes the cone of k -variate polynomials (represented by their coefficients in the standard monomial basis) nonnegative over $\Delta \subseteq \mathbb{R}^k$ with total degree n . For the applications considered in the thesis it is sufficient to consider the case when the domain Δ is polyhedral, but most of the results directly generalize to the case when Δ is a compact basic semialgebraic set. Again, we will drop the superscript whenever $\Delta = \mathbb{R}^k$. All these cones of nonnegative polynomials are closed, convex cones.

A few remarks on the terminology and notation are in order. In the related algebra literature, typically sum-of-squares *forms* (that is, homogeneous polynomials) are considered, rather than general, non-homogeneous polynomials. Every key definition and theorem about sum-of-squares forms has a counterpart involving non-homogeneous polynomials of one fewer variable. However, the notation $\Sigma_{2n,k}$ is frequently used with both meanings. It is also very common in the literature (both in algebra and in optimization) to refer to the cones \mathcal{P}_{n+1}^Δ and $\mathcal{P}_{n,k}^\Delta$ as *positive polynomials* – a terminology derived from the theory of everywhere nonnegative forms, often called *positive definite* forms. We will only use the terms *nonnegative* and *strictly positive polynomials*, as they certainly cannot create any confusion.

The connection between nonnegative and sum-of-squares polynomials goes back to the 19th century, and it is also related to Hilbert’s 17th problem [Hil00], and its affirmative solution (first obtained by Artin [Art27]). Interest in sum-of-squares polynomials has raised considerably in the past twenty years, especially in the context of global optimization of polynomials, computational algebraic geometry, and its relation to copositivity. The main reason for this is that while the recognition of nonnegative polynomials is NP-hard, the cone of sum-of-squares polynomials is semidefinite representable, and hence is tractable both in the theoretical sense, and in practice, as demonstrated by readily available software packages, such as SOSTOOLS [PPSP04], which finds sum-of-squares decompositions of multivariate polynomials; YALMIP [Löf04], a simple modeling tool capable of translating sum-of-squares polynomial constraints to semidefinite programming constraints; and GloptiPoly [HL03], a program that finds global optima of polynomials. All of these are extensions to the numerical computing environment MATLAB [Mat], and they all rely on general purpose semidefinite optimization solvers, such as SeDuMi [Stu99] and SDPT3 [TTT99]. More information on computational tools for semidefinite programming is given in the Appendix.

The semidefinite representation of the cones $\Sigma_{2n,k}$ can be derived from the following theorem on (real-valued) sum-of-squares functional systems.

Proposition 2.1 ([Nes00]). *Consider a real linear space U of $\Delta \rightarrow \mathbb{R}$ functions ($\Delta \in \mathbb{R}^k$), with a finite basis $\{u_0, \dots, u_m\}$, and let us define*

$$U^2 \stackrel{\text{def}}{=} \text{span}\{u_i u_j \mid i, j = 1, \dots, m\}.$$

Let $\{v_0, \dots, v_n\}$ be a basis of U^2 . Then the cone of sum-of-squares functions

$$\Sigma = \left\{ (p_0, \dots, p_n)^T \mid \sum_{i=0}^n p_i v_i = \sum_{i=1}^N r_i^2, \quad N \text{ finite}, r_i \in U \right\} \quad (2.1)$$

is a linear image of the cone of $(m+1) \times (m+1)$ positive semidefinite real symmetric matrices.

The linear mapping in Proposition 2.1 can be explicitly constructed in polynomial time for every given $\{u_i\}$ and $\{v_j\}$, following [Nes00]. A more transparent description of this construction can be found in the tutorial [Ali06]. We do not repeat the details here, as we are going to present a considerably more general theorem, Theorem 4.5 in Section 4.3.

An analogous representation theorem can be given for the cone of *weighted-sum-of-squares* functions, that is, functions of the form $\sum_{i=1}^{\ell} w_i s_i$, where w_1, \dots, w_{ℓ} are fixed real-valued functions, called *weights*, and each s_i belong to the cone Σ defined in (2.1).

Proposition 2.2 ([Nes00]). *With the notations of Proposition 2.1, define the cone of weighted-sum-of-squares functions as*

$$\Sigma_w = \left\{ (p_0, \dots, p_n)^T \mid \sum_{i=0}^n p_i v_i = \sum_{j=1}^{\ell} w_j \sum_{i=1}^{N_j} r_{ij}^2, \quad N \text{ finite}, r_{ij} \in U^2 \right\},$$

where w_1, \dots, w_{ℓ} are arbitrary $\Delta \rightarrow \mathbb{R}$ functions (the “weights”). Then Σ_w is the Minkowski sum of ℓ cones, each a linear image of the cone of $(m+1) \times (m+1)$ positive semidefinite real symmetric matrices, $m = \dim(U)$. The ℓ linear mappings that characterize Σ_w can be explicitly constructed in polynomial time.

We will derive a generalization of this theorem, too, in Section 4.3.1. (Theorem 4.17.)

2.1 Univariate nonnegative polynomials

It is well-known (and follows easily from the fundamental theorem of algebra) that a univariate polynomial is nonnegative if and only if it is a sum of two perfect squares; consequently, every

univariate nonnegative polynomial is sum-of-squares. This in turn yields a characterization of nonnegative polynomials as a linear image of the cone of positive semidefinite real symmetric matrices, via Proposition 2.1. These results are summarized in the following proposition. This theorem can also be derived without Proposition 2.1, from the theory of nonnegative functions in Chebyshev systems [KS66, Chapter VIII], or from the (related) theory of moment spaces [DS97].

Proposition 2.3 ([KS66, PS76]). *Suppose p is a polynomial of degree $2n$, and $p(t) = \sum_{k=0}^{2n} p_k t^k$. Then the following are equivalent.*

1. $p \in \mathcal{P}_{2n+1}$.
2. $p = q^2 + r^2$ for some polynomials q and r of degree at most n .
3. $p \in \Sigma_{2n,1}$.
4. There exists a positive semidefinite real symmetric matrix $\mathbf{X} = (x_{ij})_{i,j=0}^n$ satisfying

$$p_k = \sum_{i+j=k} x_{ij}, \quad k = 0, \dots, 2n.$$

Proposition 2.3 can be extended to the case of nonnegative polynomials over a finite or infinite interval. A natural sufficient condition for a polynomial to be nonnegative over $[a, \infty)$ is that it can be written in the form $s_1(t) + (t-a)s_2(t)$, where s_1 and s_2 are sum-of-squares polynomials. Similarly, it is sufficient for a polynomial to be nonnegative over $[a, b]$ that it is of the form $(t-a)s_1(t) + (b-t)s_2(t)$ or $s_1(t) + (t-a)(b-t)s_2(t)$, where s_1 and s_2 are sum-of-squares polynomials. It is known that these sufficient conditions are also necessary; we give an elementary proof of this fact below.

Proposition 2.4. *For every polynomial p of degree n ,*

1. $p \in \mathcal{P}_{n+1}^{[a, \infty)}$ if and only if $p(t) = r(t) + (t-a)q(t)$ identically for some sum-of-squares polynomials r and q of degree at most n ;
2. $p \in \mathcal{P}_{n+1}^{[a, b]}$ if and only if

$$p(t) = \begin{cases} r(t) + (t-a)(b-t)q(t) & (\text{if } n = 2k) \\ (t-a)r(t) + (b-t)s(t) & (\text{if } n = 2k+1) \end{cases}$$

for some sum-of-squares polynomials r and s of degree $2k$ and q of degree $2k-2$.

Proof. For simplicity we assume $a = 0$ and $b = 1$, the argument for general intervals is essentially identical. First, let us assume $p \in \mathcal{P}_{n+1}^{[0,\infty)}$, and consider the polynomial $q(t) = p(t^2)$. By assumption, q is nonnegative everywhere, hence $q = \sum_i r_i^2$. Grouping the terms of the r_i by the parity of their degrees we write $r_i(t) = r_i^{(1)}(t^2) + tr_i^{(2)}(t^2)$, and since q has no odd degree terms,

$$p(t^2) = q(t) = \sum_i \left(r_i^{(1)}(t^2) \right)^2 + t^2 \sum_i \left(r_i^{(2)}(t^2) \right)^2.$$

Taking $r = \sum_i (r_i^{(1)})^2$ and $s = \sum_i (r_i^{(2)})^2$ we have $p(t^2) = r(t^2) + t^2 s(t^2)$, implying $p(t) = r(t) + ts(t)$, as claimed. The bounds on the degrees of r and s follow from the fact that the degree of each r_i is at most n . That concludes the proof of the first claim.

Next, suppose that $p \in \mathcal{P}_{n+1}^{[0,1]}$, and consider the polynomial $q(t) = (1+t)^n p(t/(1+t))$: by assumption it is nonnegative for all $t \geq 0$, and so by the first part of the claim, $q(t) = r(t) + ts(t)$ identically for some $r = \sum_i r_i^2$ of degree at most n and $s = \sum_i s_i^2$ of degree at most $n-1$.

Observe that $p(t) = (1-t)^n q(t/(1-t))$. If $n = 2k+1$, then this yields

$$\begin{aligned} p(t) &= (1-t)^{2k+1} \sum_i r_i^2(t/(1-t)) + (1-t)^{2k+1} \frac{t}{1-t} \sum_i s_i^2(t/(1-t)) \\ &= (1-t) \sum_i \left((1-t)^k r_i(t/(1-t)) \right)^2 + t \sum_i \left((1-t)^k s_i(t/(1-t)) \right)^2. \end{aligned}$$

If $n = 2k$, then

$$\begin{aligned} p(t) &= (1-t)^{2k} \sum_i r_i^2(t/(1-t)) + (1-t)^{2k} \frac{t}{1-t} \sum_i s_i^2(t/(1-t)) \\ &= \sum_i \left((1-t)^k r_i(t/(1-t)) \right)^2 + t(1-t) \sum_i \left((1-t)^{k-1} s_i(t/(1-t)) \right)^2. \end{aligned}$$

The degree bounds come from the degree bounds of the first part of the theorem. \square

In fact we can take the polynomials q , r , and s above to be *squares* rather than sums-of-squares; this observation is usually attributed to Lukács [Luk18], and sometimes to Markov. The corresponding semidefinite characterizations can be derived using Proposition 2.2. We split these results into a sequence of propositions for better readability.

Proposition 2.5 ([Luk18, BS62]). *For every polynomial p of degree n ,*

1. $p \in \mathcal{P}_{n+1}^{[a,\infty)}$ if and only if $p(t) = r^2(t) + (t-a)s^2(t)$ identically for some polynomials r and s of degree at most $\lfloor n/2 \rfloor$;

2. $p \in \mathcal{P}_{n+1}^{[a,b]}$ if and only if

$$p(t) = \begin{cases} r^2(t) + (t-a)(b-t)q^2(t) & (\text{if } n = 2k) \\ (t-a)r^2(t) + (b-t)s^2(t) & (\text{if } n = 2k+1) \end{cases}$$

for some polynomials r and s of degree k and q of degree $k-1$.

As in the case of everywhere nonnegative polynomials, this result can also be derived independently, using the theory of moments and Chebyshev systems (see Chapter II in [KS66]), but these are rather deep tools compared to the elementary nature of these claims. See [BS62] for the first (and, to the author's knowledge, only) elementary proof, using a reduction to Proposition 2.3.

The advantage of Proposition 2.4 and its above proof is that they can be repeated essentially verbatim in every setting when a pointwise defined notion of “nonnegativity” of polynomials over the reals is known to be equivalent to the algebraic notion of “being sum-of-squares” with respect to some multiplication. (Both the more general nonnegativity and sum-of-squares notions need to possess only a few basic properties.) We shall explore this in more detail in Section 4.3, where we prove a generalization of the above theorem to univariate matrix polynomials that are positive semidefinite over an interval. We shall also use a similar technique below to derive characterizations of some nonnegative multivariate polynomials over polyhedra.

It is interesting to note that while in Proposition 2.3 it is evident that the degrees of the terms in the sum-of-squares representations of p can be bounded by the degree of p (because the highest degree terms cannot cancel each other out), this is not at all obvious in Proposition 2.5, where cancellations are possible. Having bounds on the degree of each term is critical, as Nesterov's representation theorems (Proposition 2.1 and 2.2) are only applicable to functional spaces of fixed dimension.

For the convenience of the reader seeking directly applicable recipes, we shall also present the “semidefinite programming-ready” characterization of nonnegative polynomials obtained from the application of Proposition 2.2 to the above characterization. These are instrumental in obtaining the computational results of Chapter 3.

Proposition 2.6 ([Nes00]). *Suppose p is a polynomial of degree $n = 2m + 1$, $p(t) = \sum_{k=0}^n p_k t^k$, and let $a < b$ be real numbers. Then $p \in \mathcal{P}_{n+1}^{[a,b]}$ if and only if there exist positive semidefinite real symmetric matrices $\mathbf{X} = (x_{ij})_{i,j=0}^m$ and $\mathbf{Y} = (y_{ij})_{i,j=0}^m$ satisfying*

$$p_k = \sum_{i+j=k} (-ax_{ij} + by_{ij}) + \sum_{i+j=k-1} (x_{ij} - y_{ij}) \quad (2.2)$$

for all $k = 0, \dots, 2m + 1$.

Similarly, if p is a polynomial of degree $n = 2m$, then $p \in \mathcal{P}_{n+1}^{[a,b]}$ if and only if there exist positive semidefinite real symmetric matrices $\mathbf{X} = (x_{ij})_{i,j=0}^m$ and $\mathbf{Y} = (y_{ij})_{i,j=0}^{m-1}$ satisfying

$$p_k = \sum_{i+j=k} (x_{ij} - aby_{ij}) + \sum_{i+j=k-1} (a+b)y_{ij} - \sum_{i+j=k-2} y_{ij} \quad (2.3)$$

for all $k = 0, \dots, 2m$.

It is not central to this thesis, but for the sake of completeness we should mention that the above sum-of-squares representations for nonnegative polynomials are not unique, and neither are the matrices \mathbf{X} and \mathbf{Y} in Proposition 2.6. However, Proposition 2.5 can be extended to a representation theorem that associates a *unique* representation to every polynomial.

Proposition 2.7 ([KS53, Kar63]). *Every strictly positive polynomial p of degree $2n$ admits a unique representation of the form*

$$p(t) = \alpha \prod_{i=1}^n (t - t_{2j-1})^2 + \beta \prod_{i=1}^{n-1} (t - t_{2j})^2,$$

where $\alpha > 0$, $\beta > 0$, and $t_1 < \dots < t_{n-1}$.

Every $p \in \mathcal{P}_{n+1}^{[a,b]}$ admits a unique representation of the form

$$p(t) = \begin{cases} \alpha \prod_{i=1}^m (t - t_{2j-1})^2 + \beta(t-a)(b-t) \prod_{i=1}^{m-1} (t - t_{2j})^2 & (\text{if } n = 2m) \\ \alpha(t-a) \prod_{i=1}^m (t - t_{2j-1})^2 + \beta(b-t) \prod_{i=1}^{m-1} (t - t_{2j})^2 & (\text{if } n = 2m + 1) \end{cases},$$

where $\alpha > 0$, $\beta > 0$, and $t_1 \leq \dots \leq t_{n-1}$. Moreover, p is strictly positive over $[a, b]$ if and only if in the above representation we have $t_1 < \dots < t_{n-1}$.

2.2 Multivariate nonnegative polynomials

2.2.1 Global nonnegativity

The relationship between nonnegative and sum-of-squares polynomials is more involved in the multivariate case. The first important difference from the univariate setting is that not every nonnegative polynomial is sum-of-squares. In fact, there is an “infinite gap” between sum-of-squares and nonnegative polynomials. While this was already known to Hilbert (see Proposition 2.10) the first constructive proof is due to Motzkin.

Proposition 2.8 ([Mot67]). *Let m denote the Motzkin polynomial:*

$$m(x, y) \stackrel{\text{def}}{=} 1 - 3x^2y^2 + x^2y^4 + x^4y^2.$$

The polynomial m is nonnegative over \mathbb{R}^2 (by the inequality between the arithmetic and geometric means), yet, $m + C$ is not sum-of-squares for any $C \in \mathbb{R}$.

It is unlikely that a good characterization of nonnegative multivariate polynomials exists:

Proposition 2.9. *Deciding whether a d -variate polynomial is nonnegative over \mathbb{R}^d (equivalently, minimizing a d -variate polynomial) is NP-hard even for polynomials of degree four.*

Proof. Both claims are “folklore” in the literature, but we give a short proof to underline the important equivalence of polynomial minimization and optimization with nonnegative polynomial constraints. This equivalence is shown by the following observation: for every d -variate polynomial p of total degree n ,

$$\min_{\mathbf{x} \in \mathbb{R}^d} p(\mathbf{x}) = \max_{\substack{z \in \mathbb{R} \\ p - z \in \mathcal{P}_{n,d}}} z.$$

Perhaps the easiest way to prove the first claim is by reduction from the PARTITION problem: the multiset of integers $\{a_1, \dots, a_d\}$ can be partitioned into two subsets with equal sums if and only if the minimum of the polynomial $(\sum_i a_i x_i)^2 + \sum_i (1 - x_i^2)^2$ is zero. \square

As the above proof suggests, global optimization of polynomials can be reduced to the problem of root-finding. This is a very well studied problem in algebra, and it can be solved by numerous algebraic techniques; these include algorithms using discriminants and resultants [CLO04, Chapter 7], Gröbner bases [CLO04, Chapter 2], and also homotopy (continuation) methods [Li97]. The most successful methods in practice, however, are the ones using a hierarchy of semidefinite optimization models. In other words, rather than reducing polynomial minimization to root-finding, root-finding via polynomial minimization has appeared to be successful. See also [PS03] for a brief overview of a few algebraic methods, and a computational experiment comparing various algebraic methods to a semidefinite programming technique.

The previous proof is based on a construction involving k -variate polynomials. Hence, recognizing nonnegative polynomials with a fixed number of variables and a fixed degree might still be a polynomial-time solvable problem. (In fact, it is, but the known polynomial time algorithms involve quantifier elimination techniques [Tar51, Ren92], and are neither practical, nor suitable for

optimization.) The question whether all nonnegative polynomials of a fixed degree and number of variables are sum-of-squares was answered by Hilbert.

Proposition 2.10 ([Hil88]). *The equation $\mathcal{P}_{2n,k} = \Sigma_{2n,k}$ holds in precisely three cases:*

1. *when $k = 1$, but $2n$ is arbitrary, that is, for univariate polynomials;*
2. *when $2n = 2$, but k is arbitrary, that is, for multivariate quadratic polynomials;*
3. *when $2n = 4$ and $k = 2$, that is, for biquartic polynomials.*

In all other cases, $\Sigma_{2n,k} \subsetneq \mathcal{P}_{2n,k}$.

2.2.2 Nonnegative polynomials on polyhedra

For our spline estimation problems we need characterizations of polynomials of a few variables, and of relatively small degree, that are nonnegative over simple polyhedra, such as a d -dimensional rectangle or a simplex. Without a restriction on the number of variables this is also a well-known NP-hard problem; see for example [BH02] for very simple formulations of NP-hard combinatorial optimization problems as multilinear optimization problems over the unit cube.

Proposition 2.11 ([BH02]). *Deciding whether a k -variate polynomial is nonnegative over $[0, 1]^k$ (equivalently, minimizing a k -variate polynomial over the unit cube) is NP-hard, even for multilinear polynomials of degree two.*

Nevertheless, there are several ways to approximate nonnegative polynomials over polyhedra (or even over semialgebraic sets, that is, sets defined by a finite set of polynomial inequalities) using sum-of-squares polynomials. One possibility is to represent nonnegative polynomials as sums of squares of *rational functions* whose denominators are positive over the domain of our interest. (Whether every nonnegative polynomial can be expressed as a sum of squares of rational functions was Hilbert's 17th problem [Hil00], settled in the affirmative by Artin [Art27].) This is a topic that merits its own book (indeed, see [Mar08] for a recent one), we must opt for giving only a flavor of these results, and giving pointers to the literature for the interested reader.

First let us recall Pólya's classic representation theorem of positive forms (homogeneous polynomials) on the first orthant.

Proposition 2.12 ([HLP34, Section 2.24]). *If a form is strictly positive on $\mathbb{R}_+^n \setminus \{\mathbf{0}\}$, then it can be expressed as $F = G/H$, where G and H are forms with nonnegative coefficients in the monomial basis. In particular we may suppose that $H = (\sum_{i=1}^n x_i)^r$ for a suitable r .*

With this proposition we obtain a very simple hierarchy of polyhedral inner approximations of $\mathcal{P}_{n,k}^{\mathbb{R}_+^n}$. Denoting by $F(p)$ the homogenization of p , for every $r = 0, 1, \dots$ the r -th level of the hierarchy is the set of polynomials p for which the form $(\sum_i x_i)^r F(p)$ has only nonnegative coefficients. For a stronger (semidefinite) hierarchy we may require a weaker condition: that $(\sum_i x_i)^r F(p)$ is sum-of-squares. This approach was suggested in the thesis [Par00].

As a consequence of Proposition 2.12 we obtain a representation of strictly positive polynomials over the unit simplex.

Proposition 2.13 ([HLP34, Section 2.24]). *If a polynomial is strictly positive on the simplex $\{\mathbf{x} \geq 0 \mid \sum_i x_i \leq 1\}$, then it can be expressed in the form $p(\mathbf{x}) = \sum c_{d_1, \dots, d_{n+1}} x_1^{d_1} \cdots x_n^{d_n} (1 - \sum_i x_i)^{d_{n+1}}$, where the exponents d_j are nonnegative integers, and the coefficients $c_{d_1, \dots, d_{n+1}}$ are all positive.*

For polyhedral approximations of the univariate polynomial cones the following consequence of Proposition 2.13 is also interesting.

Proposition 2.14. *Every strictly positive polynomial over $[0, 1]$ has nonnegative coefficients in the (scaled Bernstein polynomial) basis $\{x^i(1-x)^{d-i} \mid i = 0, \dots, d\}$ for a sufficiently high d .*

Nonnegativity over the first orthant and over the unit simplex are also closely related to the concept of *copositivity*. (A matrix \mathbf{A} is copositive if $\mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0$ for every $\mathbf{x} \geq 0$.) The cone of copositive matrices is very well studied, but their recognition is NP-hard, and remains a challenge even for matrices of order as small as 20. See [HUS10] for a recent review on several aspects of copositivity, and the references therein.

Proposition 2.13 (along with our univariate starting point, Proposition 2.4) motivates the other popular approach in the literature, which is also applicable to polynomials nonnegative over semialgebraic sets. If $\Delta = \{\mathbf{x} \mid w_i(\mathbf{x}) \geq 0, i = 1, \dots, m\}$, where w_1, \dots, w_m are polynomials, then a sufficient condition for a polynomial p to be nonnegative over Δ is for it to be expressible as

$$p(\mathbf{x}) = \sum_{I \subseteq \{1, \dots, m\}} \left(\prod_{i \in I} w_i(\mathbf{x}) \right) s_I(\mathbf{x}), \quad (2.4)$$

or simply as

$$p(\mathbf{x}) = \sum_{i=1}^m w_i(\mathbf{x}) s_i(\mathbf{x}), \quad (2.5)$$

where the polynomials s are sum-of-squares polynomials. The theoretical justification behind this approach is provided by theorems of Schmüdgen [Sch91] and Putinar [Put93]. Both of these theorems are generally referred to as the *Positivstellensatz* (after Hilbert's Nullstellensatz), and they show

equivalence between strictly positive polynomials on semialgebraic sets and polynomials of the form (2.4) (Schmüdgen) and (2.5) (Putinar) under different conditions. (We omit the details.) Note however, that unlike Proposition 2.5, the *Positivstellensätze* do not give explicit bounds on the degree of the sum-of-squares polynomials involved, hence these theorems are not directly useful. For more on different variants of the Positivstellensatz, as well as for bounds on the degrees of the terms required on the right-hand side of (2.4) and (2.5), see [Sch04] and [NS07]. Again, hierarchies of weighted-sum-of-squares inner approximations can be defined using higher and higher degree sum-of-squares polynomials in Equations (2.4) and (2.5), see for example [Las01b, Par03].

As a side note we shall mention that for simplexes the hierarchies built using the *Positivstellensätze* of Schmüdgen and Putinar turn out to be equivalent [KLP05] to the hierarchy based on Pólya's results above.

In our spline estimation problems we will formulate optimization models with a large number of polynomial nonnegativity constraints, so we cannot use high degree sum-of-squares polynomials required by the above techniques. In the rest of this section we shall look into the approximation of low degree nonnegative polynomials over simple polyhedra using low degree sum-of-squares polynomials. Results in this direction were obtained by Micchelli and Pinkus [MP89]. These can be viewed as the simplest successful multivariate generalizations of Proposition 2.5.

Proposition 2.15 ([MP89]). *Let $\Delta \in \mathbb{R}^2$ be any of the following polyhedra: the right half-plane $\mathbb{R}_+ \times \mathbb{R}$, the first orthant \mathbb{R}_+^2 , the unit simplex $\{(x_1, x_2) \in \mathbb{R}_+^2 \mid x_1 + x_2 \leq 1\}$, the unit square $[0, 1]^2$, or the vertical slab $\{(x_1, x_2) \in \mathbb{R}_+^2 \mid x_1 \leq 1\}$. Then every bivariate quadratic polynomial p that is nonnegative over Δ can be represented in the form (2.4) with $w_i \geq 0$ being the facet defining linear inequalities for Δ . Moreover, we may assume that there are only quadratic terms in this representation.*

The proofs of Micchelli and Pinkus involve a rather complicated case analysis and rely on the theory of copositive matrices, specifically on Diananda's result [Dia62] that every copositive matrix of order $n \leq 4$ can be written as a sum of a positive semidefinite matrix and an entrywise nonnegative matrix. This is known not to hold for $n \geq 5$, preventing the generalization of Proposition 2.15 to polynomials with more variables or of higher degree.

Most of these results can be easily derived using essentially the same technique as in the proof of Proposition 2.4. With the appropriate change of variables the nonnegativity of a polynomial p over a polyhedron is equivalent to the global nonnegativity of another polynomial q . If q belongs to one of those cones where nonnegativity is equivalent to being sum-of-squares (recall Proposition 2.10), then the sum-of-squares characterization of q can be translated back to a weighted-sum-of-squares

characterization of p . For example, to obtain Proposition 2.15 for the unit simplex, observe that the quadratic polynomial $p(x_1, x_2)$ is nonnegative over $\left\{\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \in \mathbb{R}_+^2 \mid x_1 + x_2 \leq 1\right\}$ if and only if $q(s, t) = (s^2 + t^2)^2 p\left(\frac{s^2}{s^2 + t^2}, \frac{t^2}{s^2 + t^2}\right)$ is nonnegative everywhere. (Note that this is a linear transformation $p \rightarrow q$.) But q is a bivariate quartic polynomial, hence it is nonnegative if and only if it is sum-of-squares. The characterization of those domains Δ for which this approach works is an interesting open problem.

2.3 Dual cones of nonnegative polynomials: moment cones

The dual cones of the cones of nonnegative polynomials over Δ are the *moment cones*, which are cones spanned by the moment sequences of probability measures concentrated on Δ . We shall concentrate only on the one-dimensional case here, as we will not use any multivariate results in this thesis. There is a vast literature on problems involving the representations of moments of a measure, as well as the structure of moment cones and spaces, as they have numerous theoretical and practical applications beyond optimization. See the monograph [Lan87] for a comprehensive review of applications, and [KN77, DS97] for further mathematical background on moment theory.

We will not need the “moment sequence” interpretation of the vectors in the moment cones, instead we use two other characterizations of the moment cones, these are presented in the rest of this section.

It is easy to derive a characterization of $(\mathcal{P}_{n+1}^\Delta)^*$ directly from the definition of \mathcal{P}_{n+1}^Δ . This is because $\mathbf{p} \in \mathcal{P}_{n+1}^\Delta$, that is $\sum_{i=0}^n p_i t^i \geq 0$ for every $t \in \Delta$, if and only if every vector of the form $\mathbf{c}(t) = (1, t, \dots, t^n)^\top$ satisfies $\langle \mathbf{p}, \mathbf{c}(t) \rangle \geq 0$. Hence, if we define

$$\mathcal{M}_{n+1}^\Delta = \text{cl cone} \left(\left\{ (1, t, \dots, t^n)^\top \mid t \in \Delta \right\} \right),$$

then we have $(\mathcal{M}_{n+1}^\Delta)^* = \mathcal{P}_{n+1}^\Delta$, from which we also get $(\mathcal{P}_{n+1}^\Delta)^* = \mathcal{M}_{n+1}^\Delta$. We only consider two cases: when $\Delta = [a, b]$ and when $\Delta = \mathbb{R}$; we drop the superscript Δ when $\Delta = \mathbb{R}$. We formalize these definitions below.

Definition 2.16. Using the notation

$$\mathbf{c}_{n+1}(t) \stackrel{\text{def}}{=} (1, t, t^2, \dots, t^n)^\top,$$

the *moment cone* of dimension $n + 1$ is defined as

$$\mathcal{M}_{n+1} \stackrel{\text{def}}{=} \text{cl cone}(\{\mathbf{c}_{n+1}(t) \mid t \in \mathbb{R}\}).$$

The $(n + 1)$ -dimensional *moment cone over* $[a, b]$ is defined as

$$\mathcal{M}_{n+1}^{[a,b]} \stackrel{\text{def}}{=} \text{cone}(\{\mathbf{c}_{n+1}(t) \mid t \in [a, b]\}).$$

It is not hard to show that $\mathcal{M}_{n+1}^{[a,b]}$ is closed for every $-\infty < a < b < \infty$, but to make \mathcal{M}_{n+1} closed, it is necessary to use the closure operation in the definition. More precisely,

$$\mathcal{M}_{n+1} = \text{cone}(\{\mathbf{c}_{n+1}(t) \mid t \in \mathbb{R}\}) \cup \{(0, \dots, 0, t)^T \mid t > 0\}.$$

Our first proposition is the duality relationship between these cones and the cones of nonnegative polynomials.

Proposition 2.17. $\mathcal{P}_{2n+1}^* = \mathcal{M}_{2n+1}$. Similarly, $(\mathcal{P}_{n+1}^{[a,b]})^* = \mathcal{M}_{n+1}^{[a,b]}$.

As a corollary, we can derive a semidefinite representation of moment cones from Propositions 2.3 and 2.5, using the fact that the semidefinite cone is self-dual (Proposition A.8). This gives a second characterization of moment cones; which can also be derived from the moment interpretation of these cones [ST43]. We summarize these properties in our last two propositions, which make use of the following notation.

For an odd-dimensional vector $\mathbf{c} = (c_0, \dots, c_{2m})^T$ we define its corresponding *Hankel matrix* as

$$H(\mathbf{c}) = (c_{i+j})_{i,j=0,\dots,m}. \tag{2.6}$$

Using this notation, for a vector $\mathbf{c} = (c_0, \dots, c_n)^T$ and real numbers $a < b$ we define the Hankel matrices $\bar{H}_n^{[a,b]}(\mathbf{c})$ and $H_n^{[a,b]}(\mathbf{c})$ as follows:

$$\begin{aligned} H_{2m}^{[a,b]}(\mathbf{c}) &= H(\mathbf{c}), \\ \bar{H}_{2m}^{[a,b]}(\mathbf{c}) &= H((-abc_i + (b-a)c_{i+1} - c_{i+2})_{i=0}^{2m-2}), \\ H_{2m+1}^{[a,b]}(\mathbf{c}) &= H((-ac_i + c_{i+1})_{i=0}^{2m}), \\ \bar{H}_{2m+1}^{[a,b]}(\mathbf{c}) &= H((bc_i - c_{i+1})_{i=0}^{2m}). \end{aligned} \tag{2.7}$$

Proposition 2.18 ([ST43, KS66, DS97]). *For every vector $\mathbf{c} = (c_0, \dots, c_n)^T$, $\mathbf{c} \in \mathcal{M}_{2n+1}$ if and only if the Hankel matrix $H(\mathbf{c})$, defined in (2.6), is positive semidefinite. Furthermore, with $a < b$ real numbers, $\mathbf{c} \in \mathcal{M}_{n+1}^{[a,b]}$ if and only if the Hankel matrices $\underline{H}_n^{[a,b]}(\mathbf{c})$ and $\bar{H}_n^{[a,b]}(\mathbf{c})$, defined in (2.7), are positive semidefinite.*

Optimality conditions of optimization problems are often expressed via complementarity between primal-dual vectors. Complementary pairs of nonnegative polynomials and moment cones, as well as characterization of the extreme rays of these cones, can be easily found, and they are given in the following Propositions.

Proposition 2.19. *If $\mathbf{p} \in \mathbb{R}^{n+1}$ is the coefficient vector of a polynomial p , and t is a real number, then $p(t) = \langle \mathbf{p}, \mathbf{c}_{n+1}(t) \rangle$. In particular, $p(t) = 0$ if and only if $\langle \mathbf{p}, \mathbf{c}_{n+1}(t) \rangle = 0$.*

Proposition 2.20 ([KS66, Sections 2.2 and 6.6]).

1. *The nonzero extreme vectors of \mathcal{M}_{2n+1} are the vectors $\alpha \mathbf{c}_{2n+1}(t)$ for every $\alpha > 0$ and $t \in \mathbb{R}$, and the vectors $(0, \dots, 0, \alpha)^T$ for every $\alpha > 0$.*
2. *The nonzero extreme vectors of $\mathcal{M}_{n+1}^{[a,b]}$ are the vectors $\alpha \mathbf{c}_{n+1}(t)$ for every $\alpha > 0$ and $t \in [a, b]$.*

Chapter 3

Univariate function estimation with shape constraints

Our goal in this chapter is to demonstrate applicability of optimization over nonnegative polynomials to a wide range of statistical estimation problems, specifically those with a number of shape constraints on the unknown function. Our approach is flexible enough to handle any number of additional linear constraints, such as interpolation, betweenness, or bounds on the integral of the estimator, and multiple shape constraints over different intervals. While a lot of attention has been given to shape constrained optimization, and also to employing (convex) optimization models in statistical estimation, no systematic study has appeared that demonstrate both the theoretical soundness and computational efficiency of modern convex optimization in shape-constrained estimation in such a general setting.

Much of the development of this chapter generalizes to multivariate estimation as well. However, in the multivariate setting there are a number of additional challenges; we shall discuss these separately, in Chapter 4.

In each of the estimation problems considered in this chapter the goal is to reconstruct a univariate, real-valued function f from finitely many observations (function values observed with noise, realizations of random variables, arrival times, etc.) under a collection of constraints on the shape of the function. Examples of such shape constraints include: (1) f is nonnegative over $[a, b]$, or more generally, its graph lies in a specific bounded region (defined, for example, by linear inequalities); (2) f is monotone non-decreasing or non-increasing; (3) f is convex or concave. The function f is otherwise assumed to belong to some (possibly infinite-dimensional) functional space. The optimal

estimator shall minimize a given *loss function* over the set of shape constrained functions from this space. We propose a variant of the classical *method of sieves* to find the estimator via the solution of finite-dimensional optimization problems. As our primary example we use sieves involving *polynomial splines*.

3.1 Existing results in univariate shape-constrained estimation

There is a vast literature on shape constrained estimation and learning problems. In [DTA00] a survey of smoothing regression problems with shape constraints is presented. The thesis [Mey96] reviews algorithms for both shape constrained regression and density estimation. The text of Robertson *et al.* [RWD88] gives another comprehensive survey of order restricted estimation problems, with over 800 references; Turlach [Tur05] has a more recent review on spline smoothing.

There has also been substantial research on non-parametric estimation via optimization. The works of de Montricher, Scott, Tapia and Thompson [MTT75, Sco76, STT80, TT90], and Nemirovskii, Polyak, and Tsybakov [NPT84, NPT85] are representative; in these references detailed analyses of consistency and rate of convergence of estimators are also presented. However, the majority of algorithmic techniques used to date are relatively simple and ad-hoc; often specific to the (single) shape constraint involved in the problem, with little potential for generalization. Below we summarize some of the shortcomings of the existing approaches that we would like to address in this thesis.

1. **Limited scope.** Techniques that do not explicitly involve optimization are designed to solve a specific problem, usually involving a single constraint on the shape of the estimator.

Many of the optimization-based methods also have a rather narrow scope; for example, several approaches have been proposed for shape constrained spline smoothing using splines of a specific degree. In [Hil54], convex regression using residual sum of squares loss function is approached by searching over piecewise linear functions with fixed knot points; [Bru58] employs a similar approach for monotone regression, but he works with piecewise constant functions. (In both of these papers the knot points are aligned with the data points.) In [HS98] quadratic B-splines are used for monotone regression, while [Die80] employs cubic splines in convex/concave regression—both by forcing the piecewise linear second derivative to be nonnegative at its knot points. These methods exploit the fact that global nonnegativity is expressible by finitely many linear inequalities, which will not hold even if the order of smoothness (and the degree) of the splines is increased by one. Similarly, [Tur05] proposes a method for spline smoothing that is

only applicable to quadratic and cubic splines, but not to splines of higher degree.

Whenever the optimal estimators are known *a priori* to be piecewise constant or piecewise linear, this approach is justified, and may not even require optimization. For instance, in monotone regression with residual sum of squares loss function the optimal estimator is the slope of the greatest convex minorant of the cumulative sum of the observed function values $y_i = f(x_i)$. This allowed Brunk to derive a closed form formula for the optimal piecewise constant solution [Bru58]. A similar result holds for monotone increasing or decreasing density estimation, see for example [EL01]. These ideas do not seem to generalize if we require that the estimator belong to some space of smooth functions.

The collection of papers in [Sch00] survey and justify why smoothing may be necessary for regression problems, and Eggermont and LaRiccia [EL01] make a similar case for density estimation. Hence, it is preferable to develop methods that can handle various loss functions (such as those used in penalized maximum likelihood estimation), rather than committing to a specific one.

2. **Nonlinear transformation through change of variables.** Many of the shape constraints reduce to the nonnegativity of some unknown function over its entire domain. If f is required to be nonnegative, then a simple way of imposing this constraint is to write $f = g^2$, or $f = \exp(g)$, and then turn the attention to g ; see, for example, [GG71] for a method employing such transformations. A difficulty with this approach is that linear constraints on f (such as periodicity, $f(t) = f(t + a)$, or the requirement that f integrates to one) are transformed into non-linear, and mostly non-convex constraints on g . As a result, the transformed optimization problems may become computationally intractable. A further drawback of this approach is that the underlying space of g may not be the same as that of f ; in particular, g may not belong to a finite-dimensional linear space even if f does. See [MTT75] for a study on the existence and uniqueness of maximum likelihood estimators in models involving such transformations.
3. **Overly restricted models.** At the time of some of the earlier works on shape constrained estimation, semidefinite and second-order cone programming methods were either not available or known, or were considered computationally too expensive, hence these studies often used oversimplified inexact optimization models. Most of the models reviewed and proposed in the surveys [RWD88] and [DTA00], as well as in the theses [Mey96] and [d'A04] formulate shape constraints employing only polyhedral constraints, confining themselves to proper subcones of the functional cones of interest. To see just one example of such an approach, suppose that

we are interested in nonnegative d -degree polynomial splines over an interval $[a, b]$. Observing that the B-spline basis functions are nonnegative functions that generate all polynomial splines [Boo01], it has been proposed to look at the cone generated by B-spline basis functions, which is a polyhedral cone. However, this cone is a proper subset of the cone of all nonnegative d -degree splines. The idea of I-splines [Ram88] is based on the same concept. (I-splines are integral functions of B-splines, and they form a basis of monotone nondecreasing functions.)

Since nonnegative splines can be effectively handled by modern convex optimization methods, this restriction is no longer necessary. Nonetheless, such proper subcones may in some cases provide an adequate estimator; in other cases, especially in the multivariate case, estimation using proper subcones of nonnegative may be unavoidable on complexity grounds. However, some theoretical justification must be provided to ensure that the polyhedral approximations of the cone of shape-constrained functions are adequate. We explore this question, and compare simplified models (using proper subcones) to models that use the entire cone of nonnegative functions. Specifically, we examine estimating nonnegative polynomial splines generated by piecewise Bernstein polynomial spline basis rather than B-splines, and compare the results to the SDP/SOCP approach. (We will justify our choice of Bernstein basis over B-splines in Theorem 3.6.)

4. **Under-restricted models.** To obtain simple optimization models (involving finitely many linear constraints only) some statisticians impose shape restrictions only on a *finite* subset of the domain of the unknown function. For instance, [DSTA96] and [MTA99] consider polynomial splines with nonnegativity on the k -th derivative, but they impose this nonnegativity only on the knot points. Obviously, except for piecewise linear splines, these constraints do not guarantee nonnegativity on all points. The same approach is taken in [VW87].
5. **Heuristics.** A number of iterative methods were proposed which change the number of linear constraints in each iteration, until the solution of the under- or over-restricted problem becomes a feasible solution to the original problem. For example, [Tur05] proposes an iterative method for shape-constrained spline smoothing that suffers from both over- and under-restriction. Initially nonnegativity is imposed on the spline at finitely many points. Then, if this spline is still negative somewhere, the following candidates are restricted to have critical points at those local minima of the previous spline where it assumed negative values, and nonnegativity is imposed at these critical points, too. Hence, from the second iteration onwards the optimal nonnegative spline is likely to be excluded from the search.

Ad-hoc methods for multiple shape constraints include iterative transformations of an estimator to satisfy at least one of the constraints. These methods also have no theoretical justification: even if they terminate, the resulting estimators are not known to be optimal in any sense, and in some cases they do not even yield an estimator that satisfies the required constraints.

One of our goals in this thesis is to formulate shape constrained problems at the same level of generality as it is common in the classical, not shape constrained, estimation literature, and propose flexible methods also applicable in the multivariate setting, and in the presence of multiple shape constraints.

In the univariate setting most shape constraints reduce to nonnegativity. Perhaps because in most interesting function spaces the cone of nonnegative functions is not a polyhedral cone, it seems that there is a hesitation to tackle the nonnegativity (or other shape constraints) directly in these spaces. The following three quotes from well-known references are typical.

From Ramsay [Ram88]:

“Attempting to impose monotonicity on polynomials quickly becomes unpleasant...”

From Tapia and Thompson [TT90] p. 103:

“The nonnegativity constraint [in density estimation] is, in general, impossible to enforce when working with continuous densities which are not piecewise linear.”

From Mammen and Thomas-Agnan [MTA99]:

“Constrained smoothing splines with infinitely many constraints (like $m^{(r)}(x) \geq 0$) for all x are difficult to compute...”

We show that these concerns may not be all that warranted. In particular, nonnegativity is an “easy” constraint when the shape constrained problems are formulated using polynomials, trigonometric polynomials, polynomial splines of a given order, or exponential families, and the theory of nonnegative polynomials outlined in Chapter 2 is used to reformulate these problems as semidefinite optimization problems.

In recent years some authors have already applied SDP and SOCP techniques to a few shape-constrained estimation and learning. Wang and Li [WL08] used cubic splines for isotonic regression, and Fushiki *et al.* [FHT06] have used semidefinite programming constraints with log-likelihood objective function in parametric density estimation with nonnegative polynomials. Alizadeh *et al.* [AENR08] have used SOCP and SDP models to estimate the smooth arrival rate of non-homogeneous Poisson process based on observed arrival rates using cubic splines.

The work in this thesis is an extension of these results in several directions. We apply the SDP/SOCP approach to problems of regression, density estimation, and even classification mostly in univariate problems, and we also provide theoretical background for the multivariate case. Our methods are flexible enough to allow multiple shape constraints. We may easily add interpolation constraints requiring our function to go through a number of given points, and we may impose multiple shape constraints over the entire range of data, or over a subset of the domain of the unknown function.

In Section 3.2 we lay out a sieve method for shape constrained learning problems. In Section 3.3 we apply this general approach employing polynomial splines, and show how SOCP and SDP approaches can be used effectively to solve these problems. Section 3.4 lists applications to this theory to a number of shape constrained estimation problems. The corresponding numerical results are collected in Section 3.5, where we also contrast the results to approximations based on using polyhedral cones, especially those generated by Bernstein polynomials.

We should mention that our main purpose is to demonstrate wide applicability and computational feasibility of the proposed methods, and to compare our methods to some earlier works using data offered by other researches as well as simulated data. As such, some of the computational results are mostly “proofs of concept” rather than detailed case studies of each individual problem.

3.2 Sieves and cones of nonnegative functions

3.2.1 Estimation in functional Hilbert spaces

In the most general setting in the statistical literature the goal is to reconstruct an unknown real-valued function $f \in \mathbb{H}$ in some reproducing kernel Hilbert space (RKHS) \mathbb{H} , based on finitely many *observations* $\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N$ drawn from a subset $D \subseteq \mathbb{R}^k$. The function f must satisfy a number of *shape constraints*, which translate to the requirement that f belongs to a convex, closed, pointed cone $\mathcal{K} \subseteq \mathbb{H}$, possibly intersected with an affine subspace. Finally, we seek a function f that minimizes a convex *loss functional* $L(\cdot | \mathbf{z}_1, \dots, \mathbf{z}_N)$. Hence, a *shape-constrained estimation problem* is an optimization problem of the form:

$$\begin{aligned} \inf_f \quad & L(f | \mathbf{z}_1, \dots, \mathbf{z}_N) \\ \text{subject to} \quad & A_i(f | \mathbf{z}_1, \dots, \mathbf{z}_N) = b_i \quad \text{for } i = 1, \dots, m \text{ ,} \\ & f \in \mathcal{K} \end{aligned} \tag{3.1}$$

where A_1, \dots, A_m are $\mathbb{H} \rightarrow \mathbb{R}$ linear functionals, b_1, \dots, b_m are real numbers, and L and \mathcal{K} are as explained above. Constraints on the shape of f are expressed by the conic constraint $f \in \mathcal{K}$ (possibly together with some of the linear equations). For example, \mathcal{K} can be the cone of nonnegative, monotone non-decreasing, or convex functions in \mathbb{H} . Multiple shape constraints can be modeled by considering a cone \mathcal{K} that is created the intersection of basic cones. If \mathbb{H} is finite-dimensional (which happens, for instance, when parametric models, or finite-dimensional approximations of infinite-dimensional problems are considered), \mathcal{K} is often the cone of nonnegative functions in \mathbb{H} , which we denote by \mathcal{P} in this section. Many other shape constraints, especially in the univariate case, can be reduced to this case. For example, derivatives of f may be sign-constrained, ensuring monotonicity, convexity or concavity of f .

Care has to be taken to ensure that the problem (3.1) has an optimal solution; this largely depends on the underlying space \mathbb{H} and the loss function L . Hence, \mathbb{H} is chosen as appropriate in the specific problem at hand. This is a well researched area; for instance, when a smooth estimator is sought, a Hilbert-Sobolev space $\mathbb{W}_2^m([a, b])$ is an appropriate choice for a number of commonly used loss functions [KW71]. Moreover, in certain shape constrained problems it is also known that the optimal solution belongs to a specific *finite-dimensional* space. Examples include the (unconstrained and inequality constrained) smoothing [Wah90, Chapter 1], [VW87], [EL01], and certain maximum likelihood estimation problems [NPT84]. In univariate problems, these solutions are often piecewise polynomial splines.

For an example when the problem is ill-posed, consider density estimation, where $\mathbf{z}_i = X_i$ are the realizations of a random variable whose density we seek. If the negative of the likelihood (or log-likelihood) is chosen as the loss functional, and the search space is the set of nonnegative continuous density functions, then there is no optimal solution at all.

As another example, consider the least square regression problem. In this case the data consist of $\mathbf{z}_i = (y_i, \mathbf{x}_i)$ with $\mathbf{x}_i \in \text{dom } f$. Obviously, if we choose the least square loss function $L(f|\mathbf{z}_i) = \sum_i (y_i - f(\mathbf{x}_i))^2$, then in every sufficiently “rich” space \mathbb{H} we can find solutions f which interpolate the data and thus, have a loss value of zero. The problem with such an “optimal” solution is that the resulting function, aside from being too complicated, *overfits* the data, and thus will be a poor fit at new and unobserved points \mathbf{x} .

3.2.2 Sieves and finite-dimensional approximations

The term *sieve* was coined originally by Grenander [Gre81] for a sequence of subsets S_1, S_2, \dots of some metric space of functions Ω containing the unknown function; it is assumed that $\bigcup_k S_k$ is dense in Ω . The main idea is that for any kind of statistical function estimation problem, instead of minimizing a given loss functional on the entire space Ω , search in S_m for some m , which is an easier problem. As m increases it is assumed that the “complexity” of the functions under consideration also increases. This on the one hand may give rise to a better fit to the observed data, but on the other hand, may depend too much on the data, and result in overfitting. Thus, care is taken to select the appropriate S_m . The value of m is usually determined from the data by techniques discussed below. Note that the idea is almost independent of the kind of objective function used, that is the loss functional to be minimized. Thus, it can be applied to least-squares regression problems, maximum likelihood estimation, penalized maximum likelihood estimation, and so on. Note, however, that for specific loss functionals the choice of sieves may effect the consistency and the rate of convergence of the estimators [GH82, NPT84].

The presence of shape constraints sometimes has a regularizing effect, forcing the optimal solution to belong to a simple, *a priori* known, finite-dimensional space. In such cases sieve methods may not even be necessary to avoid overfitting and ill-posedness in (3.1). For instance, if \mathcal{K} in (3.1) is the cone of monotone increasing functions, then as long as the underlying space \mathbb{H} is sufficiently large, a piecewise constant function is the least square solution [Bru58]. The same claim holds for maximum likelihood estimation of monotone probability density functions. In other situations, however, the shape constraints do not remove the overfitting and the ill-posed nature of the infinite-dimensional problems. For instance, if the only shape constraint is that f must be nonnegative (\mathcal{K} is the cone of nonnegative functions), then least squares regression will still result in interpolating functions, and maximum likelihood will result in sums of Dirac distributions. Even when the shape constraints add some regularization, they may not yield a function in the functional space of interest. In the above examples the piecewise constant functions may not be acceptable if we are seeking smooth estimators.

In this chapter we follow the sieve method, but with certain restrictions both on the space Ω and the subsets S_m . First, Ω is replaced by a RKHS $\mathbb{H} \subseteq C([a, b])$ (the set of continuous functions over $[a, b]$). Second, we assume that the sequence $\{S_m\}$ is *nested*, that is $S_0 \subset \dots \subset S_m \dots \subseteq \mathbb{H}$, or *asymptotically nested*, meaning that for every m the sieve has a nested subsequence of subsets of S_m . (We define these terms precisely below.) Finally, we assume that S_m are all finite-dimensional.

Therefore, we are considering finite-dimensional approximations of problem (3.1).

There are several ways to choose the right subcone from a sieve (that is, the right value of m), these are essentially the same approaches as those commonly used to avoid overfitting. We shall mention two of them. One approach is cross-validation, in which only a subset of the data is used for estimation, and then the resulting estimated f is tested on the remainder of the data to gauge the prediction power of the estimated function. Another approach is to measure the trade-off between the accuracy and complexity of the model using S_m . Several such measures have been proposed, they are commonly referred to as *information criteria*. For more on information criteria, see [BA04]; for model selection in general the reader is referred to the monograph [BA02].

All of these approaches can be combined with *smoothing techniques*. These include a penalty functional in the loss function which discourages “overly complex” solutions.

Let us now discuss what we mean by a finite-dimensional approximation of problem (3.1).

Definition 3.1. A sequence $\{\mathcal{K}_i\}$, $\mathcal{K}_i \subseteq \mathcal{K}$, of closed, pointed and convex cones is called a *conic sieve* if each \mathcal{K}_i is finite-dimensional, and $\bigcup_i \mathcal{K}_i$ is dense in \mathcal{K} . We say that the sieve is *nested* if $\mathcal{K}_1 \subseteq \mathcal{K}_2 \subseteq \dots \subseteq \mathcal{K}_i \subseteq \dots$, and it is *asymptotically nested* if for each i the sequence has a nested infinite subsequence containing \mathcal{K}_i .

The most straightforward construction of sieves for the cone \mathcal{K} in \mathbb{H} is to consider a countable orthonormal basis $\{\phi_i\}_{i=1,\dots}$ of \mathbb{H} , and letting $V_m \stackrel{\text{def}}{=} \text{span}(\phi_i)_{i=0}^m$, and $\mathcal{K}_i = \mathcal{K} \cap V_i$ for each i . However, it may not be computationally tractable to optimize over $\mathcal{K} \cap V_i$. As a result, we may need to find alternative conic sieves $\mathcal{K}'_i \subset \mathcal{K}_i$, where optimizing over \mathcal{K}'_i is easier than optimizing over \mathcal{K}_i . To make things concrete, let us look at a few examples.

Example: Nonnegative univariate polynomials

Suppose that in problem (3.1), $\mathbb{H} \subseteq C([a, b])$ (the set of continuous functions over $[a, b]$), and $\mathcal{K} = \mathcal{P}$ is the cone of nonnegative functions in \mathbb{H} . Then finite-dimensional approximations can be constructed using polynomials of increasing degree. Let V_i be the set of polynomials of degree at most i . Then $\bigcup_i V_i$ is dense in $C([a, b])$ (with respect to the uniform norm topology), and thus it is dense in \mathbb{H} . Then for each i , $\mathcal{K}_i = \mathcal{K} \cap V_i = \mathcal{P}_{i+1}^{[a,b]}$ is the cone of polynomials of degree i that are nonnegative over $[a, b]$. In this case, the unknown function can be expressed as $f(t) = f_0 + f_1 t + \dots + f_i t^i$, and thus, the coefficients f_j , for $j = 0, \dots, i$ are the variables of the finite-dimensional approximation. (In practical computations we would, of course, choose an orthogonal basis to represent the polynomials, not the rather poorly behaved standard basis.)

Example: Polyhedral approximations and Bernstein polynomials

If we prefer polyhedral approximations, we can use a technique similar to Ramsay's [Ram88], but for polynomials rather than splines. Namely, we should identify a basis $\{\varphi_j(t)\}_{j=0,\dots,i}$ of polynomials of degree up to i satisfying $\varphi_j(t) \geq 0$ for all $j = 0, \dots, i$ and $a \leq t \leq b$. Then we let \mathcal{K}_i be the cone generated by $\{\varphi_j(t)\}_{j=0,\dots,i}$. Care should be taken in our choice of φ_j to make sure that $\bigcup_i \mathcal{K}_i$ is dense in the cone of nonnegative functions in \mathbb{H} . For instance, the standard polynomial basis functions, $\{1, t, \dots, t^i\}$, are nonnegative over $[0, 1]$, however the cone generated by them contains only monotone non-decreasing polynomials, and cannot be used to approximate every continuous nonnegative function.

A nonnegative basis which gives rise to a conic sieve is the Bernstein polynomial basis: $B_{i,j}(t) = \binom{i}{j} t^j (1-t)^{i-j}$, $j = 0, \dots, i$. (The binomial coefficient plays no role in our context, we kept it only to be faithful to the usual definition of Bernstein polynomials.) If we choose $\mathcal{K}'_i = \left\{ \sum_{j=0}^i \alpha_j B_{i,j} \mid \alpha_j \geq 0, j = 0, \dots, i \right\}$, then $\bigcup_i \mathcal{K}'_i$ is dense in the cone of nonnegative continuous functions; this follows from the well-known proof of Weierstrass approximation theorem based on Bernstein polynomials [Lor86, Theorem 1.1.1].

Example: Univariate polynomial splines

Another important sequence of finite-dimensional function spaces is the one induced by splines. Recall Definition 1.1: a *univariate polynomial spline* f of degree d and continuity \mathcal{C}^r ($r \leq d-1$) is a real valued function on $[a, b] = [a_0, a_m]$ defined piecewise on the intervals $[a_i, a_{i+1}]$, $i = 0, \dots, m-1$ that satisfies the following properties: (1) f is a polynomial of degree d over each interval $[a_i, a_{i+1}]$, $i = 0, \dots, m-1$; (2) f has continuous derivatives up to order r over (a, b) . In this thesis we will always assume $r = d-1$. The points a_i are called the *knot points* of the spline. The linear space of all splines of order d and knot point sequence \mathbf{a} is denoted by $\mathcal{S}(d, \mathbf{a})$.

Suppose that the degree d is fixed, and V_i is chosen to be $\mathcal{S}(d, \mathbf{a}_i)$, where $\{\mathbf{a}_i\}$ is a sequence of knot point sequences such that the length of the longest subinterval in \mathbf{a}_i tends to zero. Then by [Sch81, Chapter 6] we have that $\bigcup V_i$ is dense in the Hilbert-Sobolev space $\mathbb{H} = \mathbb{W}_2^d([a, b])$, and hence $\bigcup (V_i \cap \mathcal{P})$ is dense in $\mathbb{W}_2^d([a, b]) \cap \mathcal{P}$. (As before, \mathcal{P} denotes the set of nonnegative functions in \mathbb{H} .) The sequence $\{V_i\}$ is not necessarily nested, however it is not difficult to create a subclass that forms a nested or asymptotically nested sequence. For instance we may start from one interval $[a, b]$, and then recursively add the midpoint of the rightmost longest interval to the set of knots. This way we have splines that have an almost homogeneous set of knots as the distances between adjacent knot

points are either l or $2l$ for some real number l . Similar schemes may be constructed when knots are chosen based on data \mathbf{z}_i ; in this case, of course the number of knots should grow with the number of available data. The most straightforward way to obtain an asymptotically nested sieve is to consider knot point sequences that subdivide $[a, b]$ evenly.

In this way we have a nested (or asymptotically nested) sequence of polynomial splines in $\mathbb{W}_2^d([a, b])$. Now in each spline space $\mathcal{S}(d, \mathbf{a})$ there may be a number of different convex cones that encode some type of shape constraint. Let us for the moment concentrate on the cone $\mathcal{P}^{[d, \mathbf{a}]} \stackrel{\text{def}}{=} \mathcal{S}(d, \mathbf{a}) \cap \mathcal{P}$ of nonnegative functions in $\mathcal{S}(d, \mathbf{a})$. With the notations of Chapter 2, this cone is $\mathcal{P}_{d+1}^{[a_0, a_1]} \times \dots \times \mathcal{P}_{d+1}^{[a_{m-1}, a_m]}$ intersected with an affine subspace for the required continuity of the derivatives. Therefore, $\mathcal{P}^{[d, \mathbf{a}]}$ is a convex, but in general non-polyhedral, cone.

We may wish to deal with cones of nonnegative splines that are polyhedral. Ramsey, as stated earlier, suggests working with cones generated by B-splines which form a nonnegative basis for splines with excellent numerical properties. We can also use the (polyhedral) cone of piecewise Bernstein polynomial splines. In Section 3.3.2 we prove that this cone contains the one generated by B-splines, therefore, the optimization problem with piecewise Bernstein polynomial splines in general provide lower values of loss functional. Furthermore, both cones generated by B-splines and Bernstein polynomials induce conic sieves. In the computational section of this chapter (Section 3.5), we concentrate on polynomial spline computations and compare the piecewise Bernstein polynomial cones and the full cone of nonnegative polynomial splines.

Example: Trigonometric polynomials

For many applications, in particular in signal processing, *trigonometric polynomials* may be the appropriate finite-dimensional choice. These functions are of the form $f_i(t) = \sum_{k=0}^i a_k \sin(kt) + b_k \cos(kt)$. Well-known trigonometric identities imply that the cone of nonnegative trigonometric polynomials is a nonsingular linear image of the cone of nonnegative polynomials over the entire line. (As it would take us far from our current discussion, we shall defer the proof of this claim to Section 5.1, where we also discuss some important theoretical consequences of this result.) Therefore, any representation or algorithmic method of optimization for the cone of nonnegative polynomials can easily be applied to the nonnegative trigonometric polynomials. For this reason we will not discuss trigonometric polynomials separately in this chapter.

Example: Complete Chebyshev systems

A *Chebyshev system* is a set of linearly independent univariate functions $\{u_0, u_1, \dots, u_d\}$, $u_i: [a, b] \rightarrow \mathbb{R}$, such that no linear combination of them has more than d zeros. A *complete Chebyshev system* is an infinite sequence of functions $\{u_i\}_{i=0}^\infty$, such that for each k , the set $\{u_0, u_1, \dots, u_k\}$ is a Chebyshev system. Clearly, if $V_k = \text{span}\{u_0, \dots, u_k\}$, then $\{V_k\}$ is a nested sequence of functional spaces, and the cones \mathcal{P}_k^u of nonnegative functions in V_k form a conic sieve in the functional Hilbert space in which $\{u_i\}_{i=0}^\infty$ is a Riesz basis. Karlin and Studden [KS66] thoroughly study Chebyshev and related systems, including the corresponding cones of nonnegative functions.

Sieves of monotone, convex, and concave functions

When considering sufficiently differentiable functions, many shape constraints mentioned before reduce to the nonnegativity (or nonpositivity) of the derivatives. Hence, sieves of monotone non-increasing or non-decreasing, convex or concave functions using polynomial approximations can be defined analogously to the sieves of nonnegative functions.

3.3 Representations of nonnegative polynomials and splines

3.3.1 Representations involving semidefinite and second-order cone constraints

In this section we apply the results presented in Chapter 2 to characterize nonnegative splines. We have summarized the main results on the semidefinite representation of nonnegative polynomials over an interval in Proposition 2.6. This easily extends to a characterization of nonnegative splines over $[a_0, a_m]$. As mentioned in the previous section, this immediately gives rise to analogous characterizations of non-increasing, non-decreasing, concave, and convex splines, by imposing nonnegativity on the derivatives of the spline.

An additional useful property of quadratic and cubic splines is that their characterization involve semidefinite constraints of 2×2 matrices. Positive semidefiniteness of 2×2 matrices can be translated to linear and quadratic (second-order cone) constraints using the following, well-known fact:

Proposition 3.2. *The matrix $\begin{pmatrix} x_0 & x_1 \\ x_1 & x_2 \end{pmatrix}$ is positive semidefinite if and only if $(x_0 + x_2, x_0 - x_2, 2x_1)^T \in \mathcal{Q}_3$, where $\mathcal{Q}_{n+1} = \{(x_0, \dots, x_n) : x_0 \geq \|(x_1, \dots, x_n)^T\|_2\}$ is the $(n + 1)$ -dimensional second-order cone (or Lorentz cone).*

Linear constraints and *semidefinite constraints* (that is, constraints of the form $\mathbf{X} \succcurlyeq 0$) can be handled by some nonlinear optimization and modeling software, such as CVX [GB07], SeDuMi

[Stu99], and SDPT3 [TTT03] very effectively. However, currently none of these software packages can handle an arbitrary convex nonlinear objective function, which is necessary for applications such as maximum likelihood density estimation.

Second-order cone constraints, that is, constraints of the form $\mathbf{x} \in \mathcal{Q}_{n+1}$ can be handled by most state-of-the-art convex optimization software, which also handle arbitrary convex objective functions, too. While carrying out the numerical experiments of this paper, we found the solvers KNITRO [Zie], LOQO [Van06], and CPLEX [ILO] especially effective and useful. The theoretical complexity of optimization problems involving only linear and second-order cone constraints is the same as that of linear optimization [AG03].

The equalities (2.2) and (2.3) suggest that while working with piecewise nonnegative splines, we may run into numerical problems if the interval widths (that is, the distances between adjacent knot points) are very small, or of different orders of magnitude. This can be avoided by *scaling*: apply separately an affine transformation (change of variables) on each polynomial piece $p^{(i)}$ that maps the interval $[a_i, a_{i+1}]$ (the domain of $p^{(i)}$) to $[0, 1]$, and represent each $p^{(i)}$ by the coefficients of the thus transformed polynomial, rather than by the original coefficients. The resulting *scaled representation of splines* is then the following.

$$S(x) = p^{(i)}(x) = \sum_{k=0}^n p_k^{(i)} \left(\frac{x - a_i}{a_{i+1} - a_i} \right)^k \quad \forall x \in [a_i, a_{i+1}]. \quad (3.2)$$

Analogously, we can use any other basis $U = \{u_0, \dots, u_n\}$ of polynomials of degree n in place of the monomial basis used above. The scaled representation of splines using basis U is then

$$S(x) = p^{(i)}(x) = \sum_{k=0}^n p_k^{(i)} u_k \left(\frac{x - a_i}{a_{i+1} - a_i} \right) \quad \forall x \in [a_i, a_{i+1}]. \quad (3.3)$$

For example, here is the complete list of constraints that characterize a nonnegative cubic spline of continuity \mathcal{C}^2 , with knot points a_0, \dots, a_m , with its pieces represented in the monomial basis:

Theorem 3.3. *The coefficients $p_k^{(i)}$, $i = 0, \dots, m - 1$, $k = 0, \dots, 3$ in (3.2) represent a nonnegative cubic spline over $[a_0, a_m]$ if and only if there exist real numbers $x_\ell^{(i)}, y_\ell^{(i)}$, $i = 0, \dots, m - 1$, $\ell = 0, 1, 2$ satisfying the following constraints for all $i = 0, \dots, m - 1$ (up to only $m - 2$ for the last three constraints).*

$$p_0^{(i)} = y_0^{(i)} \quad (3.4a)$$

$$p_1^{(i)} = 2y_1^{(i)} + x_0^{(i)} - y_0^{(i)} \quad (3.4b)$$

$$p_2^{(i)} = y_2^{(i)} + 2x_1^{(i)} - 2y_1^{(i)} \quad (3.4c)$$

$$p_3^{(i)} = x_2^{(i)} - y_2^{(i)} \quad (3.4d)$$

$$(x_0^{(i)} + x_2^{(i)}, x_0^{(i)} - x_2^{(i)}, 2x_1^{(i)})^T \in \mathcal{Q}_3 \quad (3.4e)$$

$$(y_0^{(i)} + y_2^{(i)}, y_0^{(i)} - y_2^{(i)}, 2y_1^{(i)})^T \in \mathcal{Q}_3 \quad (3.4f)$$

$$p_0^{(i+1)} = \sum_{j=0}^3 p_j^{(i)} \quad (3.4g)$$

$$\frac{1}{a_{i+2} - a_{i+1}} p_1^{(i+1)} = \sum_{j=1}^3 \frac{j}{a_{i+1} - a_i} p_j^{(i)} \quad (3.4h)$$

$$\frac{2}{(a_{i+2} - a_{i+1})^2} p_2^{(i+1)} = \sum_{j=2}^3 \frac{j(j-1)}{(a_{i+1} - a_i)^2} p_j^{(i)} \quad (3.4i)$$

3.3.2 Sieves of polyhedral spline cones

In this section we examine sieves comprised of polyhedral approximations. As before, we concentrate on polynomial approximations of nonnegative functions. Even though we gave a complete characterization of nonnegative splines suitable for tractable optimization models, there are still a number of reasons to consider polyhedral approximations – even in the univariate case.

Certain *combinations* of shape constraints admit a polyhedral representation, especially when splines of low degree are used. In these cases the results of the previous section are unnecessary. Consider, for example, a twice differentiable concave and nonnegative function f over $[a, b]$. Because f is concave, $f(x) \geq 0$ for every $x \in [a, b]$ if and only if $\min\{f(a), f(b)\} \geq 0$; this can be expressed by two linear inequalities. If f is approximated by a cubic spline, then the concavity constraint $f''(x) \geq 0$ ($x \in [a, b]$) also simplifies to a finite number of linear inequalities. Note, however, that nonlinear constraints are unavoidable to describe *nonnegative convex cubic* polynomials (and splines), as that cone is not polyhedral.

We concentrate on models of the following form: we fix a basis $U = \{u_0, \dots, u_n\}$ of nonnegative polynomials of degree n , and then consider splines whose coefficients $p_k^{(i)}$ in the scaled representation (3.3) are all nonnegative. Let $\mathcal{P}(U, \mathbf{a})$ denote the set of such splines with knot point sequence \mathbf{a} . These are clearly a subset of all nonnegative splines.

Our first result is a sufficient condition for a sequence of polyhedral polynomial approximations to form a sieve. For convenience, we introduce the following notation: the length of the longest subinterval in the knot point sequence \mathbf{a} (sometimes called the *mesh size* of \mathbf{a}) is denoted by $\|\mathbf{a}\|$.

Theorem 3.4. *Consider a basis $U = \{u_0, \dots, u_n\}$ of polynomials of degree $n \geq 1$ such that each*

u_i is nonnegative over $[a, b]$, and assume that $1 \in \text{int cone}(U)$, where 1 denotes the constant one polynomial. Furthermore, let $\{\mathbf{a}_i\}$ be an asymptotically nested sequence of knot point sequences satisfying $\lim \|\mathbf{a}_i\| = 0$. Then the set $\bigcup_i \mathcal{P}(U, \mathbf{a}_i)$ is a dense subcone of $\mathcal{P} \cap C([a, b])$, the cone of nonnegative functions over $[a, b]$.

Proof. Without loss of generality we can assume that $[a, b] = [0, 1]$. We start the proof by showing that for every polynomial p of degree n , strictly positive over $[0, 1]$, there exist nonnegative constants C_i such that $p + C_i \in \mathcal{P}(U, \mathbf{a}_i)$ for every i , and $\lim C_i = 0$.

Fix i , and consider two adjacent knot points a_k and a_{k+1} from the knot point sequence \mathbf{a}_i . The polynomial p can be represented as a piecewise polynomial spline of degree n with knot point sequence \mathbf{a}_i ; its scaled representation (3.2) has $p^{(k)}(x) = p((a_{k+1} - a_k)x + a_k)$. Collecting terms in the standard basis, we have

$$p^{(k)}(x) = p((a_{k+1} - a_k)x + a_k) = p(a_k) + \sum_{i=1}^n q_i^{(k)} x^i$$

with some $q_i^{(k)} = \mathcal{O}((a_{k+1} - a_k)^i)$, $i = 1, \dots, n$. By assumption, $p(a_k) > 0$, because p is strictly positive on $[0, 1]$. All other coefficients $q_i^{(k)}$ are of order $\mathcal{O}(a_{k+1} - a_k)$. By the assumption on U , $\sum_{j=0}^n \alpha_j u_j \equiv p(a_k)$ for some positive $\alpha_0, \dots, \alpha_n$. Now, if we express $p^{(k)}$ in the basis U : $p^{(k)} = \sum p_j^{(k)} u_j$, we have that $p_j^{(k)} = \alpha_j - \delta_j^{(k)}$ with $|\delta_j^{(k)}| = \mathcal{O}(a_{k+1} - a_k)$, consequently $p^{(k)} + p(a_k) \max_j (|\delta_j^{(k)}| / \alpha_j)$ has positive coefficients in the basis U . Applying the same argument for every k , we obtain that $p + C_i \in \mathcal{P}(U, \mathbf{a}_i)$ for

$$C_i = \max_{k: a_k \in \mathbf{a}_i} (p(a_k) \max_j (|\delta_j^{(k)}| / \alpha_j)).$$

Finally, as $|\delta_j^{(k)}| = \mathcal{O}(a_{k+1} - a_k)$ and p is bounded, $C_i \rightarrow 0$ as $\|\mathbf{a}_i\| \rightarrow 0$.

The same argument can be used to prove that for every strictly positive spline over $[0, 1]$, with knot point sequence \mathbf{a} , and for every sequence $\{\mathbf{a}_i\}$ consisting of subdivisions of \mathbf{a} satisfying $\lim \|\mathbf{a}_i\| = 0$, there exist nonnegative constants C_i such that $s + C_i \in \mathcal{P}(U, \mathbf{a}_i)$ for every i , and $\lim C_i = 0$.

Consequently, $\bigcup_i \mathcal{P}(U, \mathbf{a}_i)$ is a dense subset of nonnegative splines of degree n .

Finally, let us consider an arbitrary nonnegative function $f \in C([a, b])$. By the approximation theory of splines (see for example [Sch81, Theorem 6.27]), for every n there exists a constant $M_n > 0$, depending only on n , but not on f , such that for every knot point sequence \mathbf{a} there exists a (not sign-constrained) spline $s \in \mathcal{S}(n, \mathbf{a})$ satisfying

$$\|f - s\|_\infty \leq M_n \omega_n(f, \|\mathbf{a}\|), \quad (3.5)$$

where ω_n is the n th modulus of smoothness of f in $L^\infty([a, b])$, satisfying $\lim_{t \searrow 0} \omega_n(f, t) = 0$ for every $n \geq 1$, provided that f is continuous on $[a, b]$. Let ε denote the right-hand side of (3.5). Because f is nonnegative, the spline $s' = s + \varepsilon$ is strictly positive, and it satisfies

$$\|f - s'\|_\infty \leq 2M_n \omega_n(f, \|\mathbf{a}\|). \quad (3.6)$$

Hence, there are strictly positive splines $s'_i \in \mathcal{P} \cap \mathcal{S}(n, \mathbf{a}_i)$ satisfying $\lim \|f - s'_i\|_\infty = 0$.

We already saw that if $\lim \|\mathbf{a}_i\| = 0$, and $\{\mathbf{a}_i\}$ is asymptotically nested, then the positive spline s'_i with knot point sequence \mathbf{a}_i can be approximated with arbitrarily small error by some spline in $\mathcal{S}(U, \mathbf{a}_j)$ with a sufficiently high j . By (3.6), the same holds for f . \square

As Bernstein polynomials of every given degree sum to the constant polynomial 1, we can construct, for every n , a sieve that consists of polyhedral cones of n times differentiable polynomial splines.

Corollary 3.5. *For every n , the cone of piecewise Bernstein polynomial splines of degree n is a dense polyhedral subcone of nonnegative continuous functions over $[a, b]$ consisting entirely of $n - 1$ times differentiable functions.*

Let us also remark that Theorem 3.4 has straightforward generalization to multivariate splines defined piecewise over a rectilinear grid of patches. This allows us to construct polyhedral and sum-of-squares cones of splines that are dense in the set of continuous multivariate nonnegative functions. (See Section 4.1.)

Our next observation about polyhedral sets of splines is about *B-splines*. B-splines are particularly popular in the approximation and engineering literature because of their excellent theoretical and computational properties. However, as the following theorem shows, cones generated by B-splines of a given degree and knot point sequence form a proper subcone of splines with nonnegative coefficients in their scaled representation (3.3) with the Bernstein polynomial basis as U . Hence, we do not consider B-splines any further.

Theorem 3.6. *For every positive integer n , the cone of functions generated by B-splines of degree n with knot points \mathbf{a} is a subset of the cone of piecewise Bernstein polynomial splines of the same degree, with knot points \mathbf{a} . For $n \geq 2$ this containment is strict.*

Proof. B-splines are commonly represented as a linear combination of the B-spline basis functions, whose supports are intervals between non-consecutive knot points. Recall the famous Cox–de Boor

recursion formula [Boo01]: if $f_{n,i}$, $i = 0, \dots, m$ denote the B-spline basis functions of degree n with knot points $0, \dots, m-1$, then for every $i = 0, \dots, m$,

$$f_{0,i}(x) = \begin{cases} 1 & i \leq x < i+1 \\ 0 & \text{otherwise} \end{cases}, \quad \text{and} \quad (3.7a)$$

$$f_{n,i}(x) = \frac{x-i}{n} f_{n-1,i}(x) + \frac{(i+n+1)-x}{n} f_{n-1,i+1}(x) \quad n \geq 1. \quad (3.7b)$$

To prove our claim, we need a characterization of B-splines of degree n as splines of the form $\mathcal{S}(U_n, \mathbf{a})$ with some basis U_n of degree n polynomials. First, we obtain a recursion on the appropriate U_n from (3.7), and then we establish a recursive formula on the basis transformation matrix between U_n and the degree n Bernstein polynomial basis. We complete the proof by showing that each entry of that matrix is nonnegative.

A B-spline segment between two consecutive knot points can be written as a nonnegative linear combination of the same segment of the B-spline basis functions. We can assume without loss of generality that the knot points in question are at 0 and 1, and to simplify the calculations we consider scaled Bernstein polynomials with leading coefficients ± 1 , that is, the i th Bernstein polynomial of degree n is $b_{n,i}(x) = x^i(1-x)^{n-i}$.

Note that the basis functions $f_{n,0}, \dots, f_{n,m}$ are shifted copies of each other: $f_{n,i}(x) = f_{n,0}(x-i)$. Hence the B-spline segment between knot points 0 and 1 is a nonnegative linear combination of the functions $f_{n,0}(x+i)$, $i = 0, \dots, m$, or more precisely, of the functions

$$g_{n,i}(x) \stackrel{\text{def}}{=} f_{n,0}(x+i) \quad \forall i = 0, \dots, n,$$

as $f_{n,0}(x+i)$ is identically zero for $i > n$. Hence, $U_n = \{g_{n,0}, \dots, g_{n,n}\}$. Rewriting the recursion (3.7b) in terms of $g_{n,i}$ we obtain

$$\begin{aligned} g_{n,i}(x) &= f_{n,0}(x+i) \\ &= \frac{x+i}{n} f_{n-1,0}(x+i) + \frac{(n+1)-(x+i)}{n} f_{n-1,1}(x+i) \\ &= \frac{x+i}{n} f_{n-1,0}(x+i) + \frac{(n+1)-(x+i)}{n} f_{n-1,0}(x+(i-1)) \\ &= \frac{x+i}{n} g_{n-1,i}(x) + \frac{(n+1)-(x+i)}{n} g_{n-1,i-1}(x), \quad \forall i = 0, \dots, n, \end{aligned} \quad (3.8)$$

with $g_{0,0} = 1$.

Using this recursion we can also find a recursive formula for the matrix of the basis transformation

from the polynomials $g_{n,i}$ to the Bernstein polynomials $b_{n,i}$. Let $A_{n,k,i,j}$ be the coefficient of $b_{n,j}$ in the unique representation of $g_{k,i}$ in the basis $\{b_{n,0}, \dots, b_{n,n}\}$, and $B_{n,k,i,j}$ be the coefficient of $b_{n,j}$ in the unique representation of $x \rightarrow xg_{k,i}(x)$ in the same basis. Then we have

$$A_{n+1,n,i,j} = A_{n,n,i,j} + A_{n,n,i,j-1} \quad (3.9)$$

from the identity $b_{n+1,j-1} + b_{n+1,j} = b_{n,j}$, and

$$B_{n+1,k,i,j} = A_{n,k,i,j-1} \quad (3.10)$$

from the identity $xb_{n,j-1}(x) = b_{n+1,j}(x)$. (The coefficients $A_{n,k,i,j}$ are defined to be zero whenever the indices are “out of bounds”, in particular when $j < 0$.) Using (3.9) and (3.10), the recursion (3.8) on $g_{n,i}$ translates to

$$\begin{aligned} nA_{n,n,i,j} &= B_{n,n-1,i,j} + iA_{n,n-1,i,j} + (n+1-i)A_{n,n-1,i-1,j} - B_{n,n-1,i-1,j} \\ &= A_{n-1,n-1,i,j-1} + iA_{n,n-1,i,j} + (n+1-i)A_{n,n-1,i-1,j} - A_{n-1,n-1,i-1,j-1} \\ &= A_{n-1,n-1,i,j-1} + iA_{n,n-1,i,j} + iA_{n-1,n-1,i,j-1} + (n+1-i)A_{n-1,n-1,i-1,j} + \\ &\quad + (n+1-i)A_{n-1,n-1,i-1,j-1} + A_{n-1,n-1,i-1,j-1} \\ &= iA_{n-1,n-1,i,j} + (n+1-i)A_{n-1,n-1,i-1,j} + \\ &\quad + (i+1)A_{n-1,n-1,i,j-1} + (n-i)A_{n-1,n-1,i-1,j-1} \end{aligned}$$

for every $n \geq 1$ and $0 \leq i, j \leq n$. Every coefficient on the right-hand side is nonnegative. Since $A_{0,0,0,0} = 1 \geq 0$, this implies that $A_{n,n,i,j} \geq 0$ for every $0 \leq i, j \leq n$ by induction.

To see that the containment is strict when $n \geq 2$, it is sufficient to exhibit a Bernstein polynomial that is not a nonnegative linear combination of U_n . Since $b_{n,k}(0) = b_{n,k}(1) = 0$ for every $k = 1, \dots, n-1$, but $g_{n,i}(0) + g_{n,i}(1) > 0$ for every $i = 0, \dots, n$, it follows that every $b_{n,k}$ is such a polynomial for $k = 1, \dots, n-1$. \square

3.4 Shape-constrained statistical estimation models

This section summarizes a number of univariate non-parametric statistical estimation problems where the techniques of Sections 3.2 and 3.3 are directly applicable. Numerical illustrations follow in Section 3.5.

3.4.1 Non-parametric regression of a nonnegative function

One of the most fundamental problems in statistics is univariate regression, the problem of estimating a univariate function f based on data $\mathbf{z}_i = (x_i, y_i)$ $i = 1, \dots, N$, assumed to have come from the model

$$y_i = f(x_i) + \varepsilon_i \quad \forall i,$$

where f is the unknown function to be estimated, and ε_i are independent, identically distributed random variables with mean zero.

The function f is assumed to belong to a class of functions \mathbb{H} , and is estimated by finding a function in \mathbb{H} that fits the data at least as well as any other function in \mathbb{H} . The goodness-of-fit of f to the data is measured by some loss functional of the form $L(f|\mathbf{z}_1, \dots, \mathbf{z}_N) = d(f) + s(f)$, where the term $d(f)$ measures the distance of the function values $f(x_i)$ and y_i , while $s(f)$ is a penalty term that penalizes non-smooth or in other ways overly complex solutions.

The most common choice for $d(f)$ is the *residual sum of squares* $d(f) = \sum_{i=1}^N (f(x_i) - y_i)^2$, which is popular for many reasons: it is simple, easy to handle analytically, often leads to models that can be solved by a closed form formula, and it appears in maximum likelihood models if ε_i are normally distributed.

The smoothing term $s(f)$ is often omitted, especially if \mathbb{H} already consists of smooth functions only. If present, it is typically $\int |f'|$, $\int |f''|$, or $\int (f'')^2$ [EL01].

When there are additional shape constraints on the function f , and when we are estimating f by splines of degree three or less, then all the above choices of $s(f) + d(f)$ lead to optimization models with only linear and second-order cone constraints. (Higher degree splines lead to general semidefinite constraints.)

We should note that in the case of nonnegative regression using least squares error model, or assuming that the error is normally distributed with mean zero, implies that while we are looking for a nonnegative estimator, the observations *can* be negative. If we assume that observations are also nonnegative, then the normal error model is not appropriate.

3.4.2 Isotonic, convex, and concave non-parametric regression

Suppose the estimator of the unknown function f belongs to a linear space of sufficiently differentiable functions whose derivatives lie in a space where nonnegativity is easily characterized. Then monotonicity, concavity, and convexity constraints on the estimator can be added to the optimization models without any difficulty. In particular, monotone, convex, and concave splines can be characterized.

When using cubic splines, convexity and concavity can be expressed by linear constraints.

3.4.3 Unconstrained density estimation

One can formulate the estimation of probability density functions as a shape constrained optimization problem. The problem is to estimate an unknown probability distribution from a finite set of independent samples $\{X_1, \dots, X_n\}$ of that distribution.

A PDF must be nonnegative and integrate to one. Formally, we assume that the PDF to be estimated has finite support, say $[a_0, a_m]$, and that it is continuous, therefore it can be approximated by nonnegative polynomial splines of a fixed degree. When using a spline model, the condition that the PDF should integrate to one simplifies to a linear constraint, since the integral of a polynomial on a given interval is a linear function of the coefficients of the polynomial. For example, a cubic spline model can be constructed by adding to (3.4) the constraint

$$\sum_{i=0}^{m-1} \sum_{j=0}^3 \frac{a_{i+1} - a_i}{j+1} p_j^{(i)} = 1$$

Finally, the objective function needs to be determined. The most common and straightforward approach is *maximum likelihood* estimation. If the unknown PDF is denoted by f , this amounts to maximizing the likelihood function $\prod_{i=1}^n f(X_i)$. This objective function will cause numerical problems, and it is not necessarily a concave function either. Instead, we will use the log-likelihood function $\sum_{i=1}^n \log f(X_i)$ as the objective function, which is concave if f is a polynomial spline of a given knot sequence. It is also important to note that by constraining f to be a polynomial spline, the above maximum likelihood optimization problem is always well-defined. (It has an optimal solution for any fixed set of knot points.)

Alternatively, one may wish to estimate a cumulative distribution function of a distribution supported on $[a_0, a_m]$. Then the constraints to be added are: $p(a_0) = 0$, $p(a_m) = 1$ (these are two linear constraints), and the nonnegativity of p' , another system of constraints that express the nonnegativity of a polynomial. The nonnegativity of the CDF itself is then superfluous, and the objective function has to be changed, too. We omit the details and instead focus on PDF estimation.

3.4.4 Unimodal density estimation

Further constraints can be added to these models for *unimodal* density estimation. If the mode is known, we can place one of the knot points on the mode, and then add constraints that the spline

is increasing from the first knot point to the mode, and decreasing from the mode to the last knot point.

If the mode is unknown, we have a non-convex problem: the set of unimodal functions is not convex simply because sum of two unimodal functions is not necessarily unimodal. In this situation an approximate solution to the problem can be found by solving a sequence of optimization models, each with a different mode, and comparing the optimal solutions that correspond to the different modes. Finding the exact solution this way is still nontrivial, as the maximum likelihood function (as a function of the mode) is not unimodal, let alone concave.

Owing to the unimodality constraints, the nonnegativity constraints may be simplified: it is sufficient to require the PDF to be nonnegative at the first and last knot point.

3.4.5 Log-concave density estimation

Another frequently useful shape constraint that appears in many applications is log-concavity. In this section we investigate how to apply the methods discussed in the previous sections to the maximum likelihood estimation of log-concave densities. Distributions with log-concave density include many common parametric family of continuous probability distributions, such as the normal, uniform, χ^2 distributions (for all parameter values), the beta and Gamma distributions (for some parameter values). They are known for a variety of applications, for example, in econometrics, reliability theory, and mechanism design (see [BB05] for a detailed exposition and references), as well as in stochastic programming [Pré71].

Recall that a nonnegative function f is *log-concave* (or *logarithmically concave*) if $f(x)^\lambda f(y)^{1-\lambda} \leq f(\lambda x + (1-\lambda)y)$ for every $\lambda \in [0, 1]$ and $x, y \in \text{dom } f$. Alternatively, f is log-concave if $f = \exp(g)$ for some concave function $g: \text{dom } f \rightarrow [-\infty, +\infty)$. We consider the following basic setup: given independent and identically distributed observations x_1, \dots, x_n , our goal is to find a probability density function f that maximizes the log-likelihood $\sum_{i=1}^n \log f(x_i)$ among smooth logarithmically concave density functions.

It was independently shown recently both in [PWM07] and [Ruf07] that without the smoothness requirement the maximum likelihood estimator is unique, and that its logarithm is a piecewise linear function whose knot points are the observations x_i . The R package `logcondens` implements two algorithms that find this estimator. The consistency and the rate of convergence of this estimator are established in [DR09]. A piecewise linear estimator for the logarithm, however, may be unsatisfactory if we are looking for a smooth estimator. Kernel methods with smooth kernels yield

smooth estimators, but they necessitate expensive cross-validation for the bandwidth parameter, and there is no guarantee that the resulting estimator is optimal in any sense. Moment problems with log-concavity and unimodality constraints are also well studied, see for example [Pré90, SSP09] for an optimization-based approach to the univariate discrete moment problem (with and without shape constraints), and [MNP04] for a study on the multivariate case. Discrete moment problems are outside the scope of this thesis.

The set of log-concave functions is not convex, and hence it is generally difficult to optimize over them directly. Instead, the logarithm of the density can be estimated, whose concavity is a convex constraint. All algorithms in the papers cited above use the same approach, and we shall adopt it, too.¹ In the remainder of this section we show that finding a maximum likelihood polynomial spline estimator for the logarithm of the unknown density amounts to solving a simple convex nonlinear optimization problem.

Let $f = \exp(g)$ be the probability density function over $[a, b] \subset \mathbb{R}$. Choosing g to be the variable of our optimization problem raises some difficulty: the linear constraint $\int_a^b f = 1$ becomes the nonlinear equality $\int_a^b \exp(g) = 1$. However, as the next lemma shows, the latter constraint can be changed to $\int_a^b \exp(g) \leq 1$ without loss of generality, and as we show below (see Lemma 3.8) this change leads to a convex programming formulation.

Lemma 3.7. *Let $[a, b] \subseteq \mathbb{R}$ a closed interval, and K be a cone containing only integrable functions over $[a, b]$, at least one of which is strictly positive. Given numbers $x_1, \dots, x_n \in [a, b]$, consider the optimization problem*

$$\text{maximize } \sum_{i=1}^n \log f(x_i) \tag{3.11a}$$

$$\text{subject to } \int_a^b f(t) dt \leq 1 \tag{3.11b}$$

$$f \in K \tag{3.11c}$$

If the optimal solution exists, then it satisfies (3.11b) with equality.

Proof. Let f be a feasible solution to (3.11), and $I = \int_a^b f(t) dt$. If $I = 0$, then the problem is unbounded, because λf is also feasible for every $\lambda > 0$. If $0 < I < 1$, then f cannot be optimal,

¹We realize that in this case we are using an approach that we criticized in the introductory section, namely we are employing a nonlinear transformation of the function. However, in the case of log-concavity constraint, there are no general tractable methods without this transformation. As such, the range of problems that can be solved is also limited. For instance, we cannot have many more linear equality constraints, since in the $f \rightarrow \exp(f)$ transformation linearity and convexity are lost.

because f/I is also feasible to (3.11), with higher objective function value. Hence, every optimal solution must satisfy $I = 1$. \square

In our application, K is the set of functions of the form $f = \exp(g)$, where g is a concave spline of a given degree, positive over $[a, b]$, with a given set of knot points. Problem (3.11) is now equivalent to

$$\text{maximize } \sum_{i=1}^n g(x_i) \tag{3.12a}$$

$$\text{subject to } \int_a^b \exp(g(t)) dt \leq 1 \tag{3.12b}$$

$$g \in \mathcal{P} \cap \mathcal{S}(d, \mathbf{a}) \tag{3.12c}$$

for the d and \mathbf{a} of our choice.

Since (3.12b) is a convex constraint (see below), the set of feasible solutions to this problem is convex, closed, and non-empty. It must also be bounded, because otherwise there would be a nonzero spline that is nonnegative over $[a, b]$, and which integrates to zero. Consequently (3.12) always has an optimal solution.

It only remains to show that (3.12b) is indeed a convex constraint. This is an immediate consequence of the next lemma, which we state in a slightly more general form than we need, so it can be applied to other spaces of functions, such as trigonometric polynomials.

Lemma 3.8. *Given any interval $[a, b] \in \mathbb{R}$ and componentwise integrable function $\mathbf{u}: [a, b] \rightarrow \mathbb{R}^n$, the function $F: \mathbb{R}^n \rightarrow \mathbb{R}$ defined by the formula*

$$F(\mathbf{p}) = \int_a^b e^{\langle \mathbf{p}, \mathbf{u}(t) \rangle} dt$$

is convex in \mathbf{p} .

Proof. Using Hölder's inequality and the weighted version of the inequality of arithmetic and geometric

means we obtain that for every $0 < \lambda < 1$,

$$\begin{aligned}
F(\lambda \mathbf{p}_1 + (1 - \lambda) \mathbf{p}_2) &= \int_a^b e^{\langle \lambda \mathbf{p}_1 + (1 - \lambda) \mathbf{p}_2, \mathbf{u}(t) \rangle} dt \\
&= \int_a^b e^{\lambda \langle \mathbf{p}_1, \mathbf{u}(t) \rangle} e^{(1 - \lambda) \langle \mathbf{p}_2, \mathbf{u}(t) \rangle} dt \\
&\leq \left(\int_a^b e^{\langle \mathbf{p}_1, \mathbf{u}(t) \rangle} dt \right)^\lambda \left(\int_a^b e^{\langle \mathbf{p}_2, \mathbf{u}(t) \rangle} dt \right)^{1 - \lambda} \\
&\leq \lambda \int_a^b e^{\langle \mathbf{p}_1, \mathbf{u}(t) \rangle} dt + (1 - \lambda) \int_a^b e^{\langle \mathbf{p}_2, \mathbf{u}(t) \rangle} dt \\
&= \lambda F(\mathbf{p}_1) + (1 - \lambda) F(\mathbf{p}_2). \quad \square
\end{aligned}$$

Setting the components of \mathbf{u} to be a basis of $\mathcal{S}(d, \mathbf{a})$ in Lemma 3.8 shows that the inequality (3.12b) is indeed a convex constraint.

3.4.6 Arrival rate estimation

The recent paper [AENR08] gives the details of several maximum likelihood models for estimating the arrival rate of a non-homogeneous Poisson process using nonnegative splines. The models are different primarily in their assumptions on how the data are given: by exact arrival times, or by the number of arrivals in fixed, short, intervals. Furthermore, slightly different models can be obtained if we assume that the arrival rate function is periodic with a fixed period length.

Compared to the regression model with a nonnegativity constraint, the nonnegative spline model for the arrival rate approximation problem is only different in its objective function. The least squares objective function is replaced by a likelihood (or log-likelihood) function, and minimization by maximization. The natural log-likelihood approach directly leads to a tractable model with concave objective function, the details can be found in [AENR08]. We show a multivariate version of this approach Chapter 4.

The same paper also compares the shape constrained spline approach to a number of other approaches. Hence, in Section 3.5 we only compare the SOCP-based model to the linearly constrained model (using Bernstein polynomials), on simulated data, and also on the same dataset of email arrivals as the one used in [AENR08].

3.4.7 Knot point selection

In each of the spline models above we have assumed that a fixed sequence of knot points a_0, \dots, a_m is given. Finding the best selection of knot points is a central, but very difficult, problem. Ideally, we would make the knot points variables, and optimize over them as well as the coefficients of the polynomials, but this would result in a intractable, non-convex optimization problem.

A common practice is to use evenly distributed knot points. While this is a very crude method, it is also very simple, and it results in simplified optimization models. (For example, from (3.4) we can eliminate the parameters a_i entirely.) It also fits the conic sieve framework discussed in Section 3.2, as splines with uniform knot points give rise to an asymptotically nested sieve as we increase the number of knot points. Since our primary goal is to show that our shape constrained spline models are competitive with the most commonly used kernel methods in all of the applications mentioned above, we decided to use this simple knot point selection method. As shown in Section 3.5, this was already sufficient for good results.

Another knot point selection method commonly used in regression problems is to place knot points at the data points. A theoretical result supporting this idea is that the optimal solutions to certain optimization models (least squares regression with a penalty term for smoothing) are natural cubic splines [Wah90]. Even when such a result is not available for our problem, we can start with a trivial subdivision (with two knot points, one at each endpoint of the domain), and add knot points one by one, at each step subdividing one of the intervals with the largest number of data points in it.

Finally, we must also decide on the number of knot points. We can accomplish this task by any of the common model selection procedures: validation on a test set (if there is one), cross-validation or k -folding (if there is none), or by measuring the trade-off between goodness-of-fit and the number of estimated parameters using either the Akaike or the Bayesian Information Criterion (AIC and BIC, respectively) [BA02, BA04]. For more on the details of various common cross-validation methods, see [AC10].

While cross-validation and k -folding are much more time consuming than computing a score such as AIC or BIC, all models proposed in this chapter are solvable efficiently enough that even using leave-one-out cross-validation was computationally feasible in the examples where cross-validation of parameters was necessary.

When the knot point sequences were not nested, but asymptotically nested, we used leave-one-out cross-validation. Smoothing parameters in objective functions were cross-validated similarly. Whenever we used nested sieves, and the objective function had no smoothing penalty, we used

AIC_c for model selection for simplicity.²

3.5 Numerical examples

In this section we have compiled results of a number of numerical experiments in which the SOCP-based sieves of cubic and quartic splines (nested and asymptotically nested) were compared to piecewise Bernstein polynomial spline models, and also to kernel methods, and in some cases to parametric methods, whenever these were available. The reader should keep in mind that owing to the large number of models and problems, these results are meant only as numerical illustrations, which demonstrate the wide applicability and computational feasibility of the models of Sections 3.3 and 3.4.

Acronyms are used to refer to the models in the tables. PM refers to Parametric Model, the precise meaning is explained in each section, as different parametric models are used in different problems. The acronyms for spline-based models are combinations of the letters C, Q, N, I, and V, which stand for Cubic, Quartic, Nonnegative, Increasing, and *concaVe* splines, respectively.

The optimization models described in Sections 3.3 and 3.4 were implemented using the AMPL modeling language [FGK02], and were solved using the nonlinear solvers KNITRO [Zie], version 5.1, or CPLEX [ILO], version 11.2.

3.5.1 Isotonic and convex/concave regression

We compared various shape-constrained spline models to parametric approaches using a dataset arising from applications as well as simulated data. In this, and also in each of the following, sections we will indicate the reference to the parametric model which arose from real applications, whenever available, along with the datasets used. For parametric models from simulated data we use simple one and two parameter families of functions which contain the known true function used to simulate the data.

² AIC_c is a variant of AIC with an additional second-order term that becomes necessary when the number of estimated parameters is small [BA02]. When the number of estimated parameters tends to infinity, the difference between the values of AIC_c and AIC tends to zero, so AIC_c is justified regardless of the number of parameters.

Isotonic regression with simulated data

We simulated noisy data using the model $Y = f(X) + \varepsilon$, where f is a smooth step function given by the formula

$$f(x) = 5 + \sum_{i=1}^4 \operatorname{erf}(15i(x - i/5)) \quad x \in [0, 1], \quad (3.13)$$

where $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-x^2}$ is the error function, and ε is normally distributed with mean 0. The standard deviation σ of ε was varied in different experiments. The function f was chosen so that it is increasing, yet it has a number of essentially flat sections, as well as strictly increasing sections of various slopes. As a result, this function is likely to expose the shortcomings of regression methods that do not include monotonicity as a constraint in their model. This function also exposes the oscillation problems that often arise in estimation with polynomials.

Random samples of size 100 were drawn uniformly from the interval $[0, 1]$. Optimal cubic spline estimators were selected from the unconstrained and monotone increasing cubic splines with up to 30 uniformly placed knot points. We compared the estimators by measuring the L_2 , L_1 , and L_∞ distances of the true function f and the estimators. This process was repeated 100 times with $\sigma = 0.15$ and 100 times with $\sigma = 0.3$. The means and standard deviations of the L_2 , L_1 , and L_∞ distances are reported in Tables 3.1 and 3.2. In the table headings μ_p and σ_p denote the estimated mean and standard deviation of the L_p distance; $p \in \{1, 2, \infty\}$.

We also considered smoothing splines: cubic splines that minimize the penalized residual sum of squares objective function

$$d(f) + s(f) = \sum_{i=1}^N (f(x_i) - y_i)^2 + \lambda \int_0^1 (f''(x))^2 dx,$$

where $\lambda > 0$ is the smoothing parameter. The optimal number of knot points and the optimal λ were both determined by cross-validation.

model	$10^2 \mu_2$	$10^2 \sigma_2$	μ_∞	σ_∞	$10 \mu_1$	$10 \sigma_1$
unconstrained	1.634	1.014	0.459	0.230	0.880	0.148
monotone	1.344	0.412	0.481	0.114	0.749	0.118
smoothing	1.678	0.825	0.455	0.159	0.904	0.166
monotone, smooth	0.807	0.477	0.339	0.127	0.610	0.113

Table 3.1: Comparison of the unconstrained and the monotone spline estimators, with and without smoothing penalty, for the regression curves of a dataset simulated using the model (3.13) with $\sigma = 0.15$. μ_i and σ_i are the estimated mean and standard deviation of the L_i errors based on 100 experiments.

Sample plots of some of the optimal estimators are shown in Figure 3.1. These plots are typical

model	$10^2 \mu_2$	$10^2 \sigma_2$	μ_∞	σ_∞	$10\mu_1$	$10\sigma_1$
unconstrained	4.619	7.798	0.659	0.513	1.475	0.237
monotone	4.426	0.911	0.715	0.110	1.478	0.177
smoothing	5.359	7.279	0.737	0.536	1.581	0.388
monotone, smooth	4.052	7.799	0.581	0.307	1.290	0.679

Table 3.2: Comparison of the unconstrained and the monotone spline estimators, with and without smoothing penalty, for the regression curves of a dataset simulated using the model (3.13) with $\sigma = 0.3$. μ_i and σ_i are the estimated mean and standard deviation of the L_i errors based on 100 experiments.

in that the unconstrained splines generally showed considerably more oscillation than the monotone estimators (which, of course, cannot oscillate).




Figure 3.1: Unconstrained and monotone spline estimators for the regression curve of a simulated dataset (with $\sigma = 0.3$). The function (3.13), to be estimated, is shown in black (thin line), the thick blue curves are the estimators. Top: best least-squares fits. Bottom: smoothing splines. Left: unconstrained splines. Right: monotone splines.

It is important to note that the unconstrained model did not yield a monotone estimator in any of the 200 experiments. In fact, even unconstrained smoothing splines did not yield monotone estimators. As smoothing did not yield significant improvement in these experiments (nor it helps satisfying the shape constraints), we did not consider smoothing splines in further experiments.

Mixed shape constraints – the rabbit eye–age data

This data is from [DM61], collected in a study to find the relationship between the weight of the eye lens and the age of rabbits in Australia. The 71-sample dataset was retrieved from the OzDASL data library [OzD]. The same data was used recently in [BD07] for comparison of shape-constrained

regression methods, as biological evidence suggests that the relationship between the age and the size of rabbits is both monotone and concave.

A further difficulty is that the data appears to be heteroscedastic. (The weight and size of adult rabbits presumably has more variation than those of young rabbits.) The original solution in [DM61] is based on parametric least-squares regression, after an ad-hoc transformation of the variables that reduces the heteroscedasticity of the data. The model for regression is $a = A \cdot 10^{-B/(w+C)} + \varepsilon$, where the response variable a is the age and the covariate w is the dried eye lens weight of the rabbits; A , B , and C are the unknown nonnegative parameters. This parametric model also yields nondecreasing functions for all parameter values, and concave functions whenever $B \leq 2C$. The solution reported in [DM61] was $A = 288.27$, $B = 60.1912$, $C = 41$, indeed increasing and concave; close to the best ordinary least squares fit of this model, which is $A = 279.8278$, $B = 55.4063$, $C = 36.0437$. Below we demonstrate that a shape-constrained non-parametric approach, combined with techniques to reduce the effect of heteroscedasticity, can yield a comparably good fit.

We added both shape constraints, separately and together, to the nested sieve of quartic splines. (Each added knot point subdivides the interval with the largest number of samples as evenly as possible.) To reduce the effect of heteroscedasticity, we used a weighted residual sum of squares loss function with weights determined by the iteratively re-weighted least squares heuristic of Schlossmacher [Sch73].

Plots of the estimated regression curves are shown on Figure 3.2. It is clear that both shape constraints are necessary to obtain an increasing and concave estimator, that is, both constraints are active at the solution of the optimization problem.

Mixed shape constraints – simulated data

We simulated noisy data using the model $Y = f(X) + \varepsilon$, where $f(x) = \frac{1}{1+e^{-10x}}$, $x \in [0, 1]$, and ε is normally distributed with mean 0 and standard deviation 0.2. This function was chosen so that the function has a nearly linear, increasing, and a long, nearly horizontal part on the domain – this way it is likely that explicit monotonicity and concavity constraints will be required for a good quality fit.

Random samples of size 50 were drawn uniformly from the interval $[0, 1]$. For each sample the least-squares optimal model from the one-parameter family $f_b(x) = \frac{1}{1+e^{-bx}} + \varepsilon$ and the two-parameter family $f_{a,b}(x) = \frac{1}{a+e^{-bx}} + \varepsilon$ were computed. In this section we refer to these reference models as parametric models PM1 and PM2. Since these least-squares problems are non-convex optimization problems, we verified (by brute-force) that the locally optimal solutions found are indeed globally

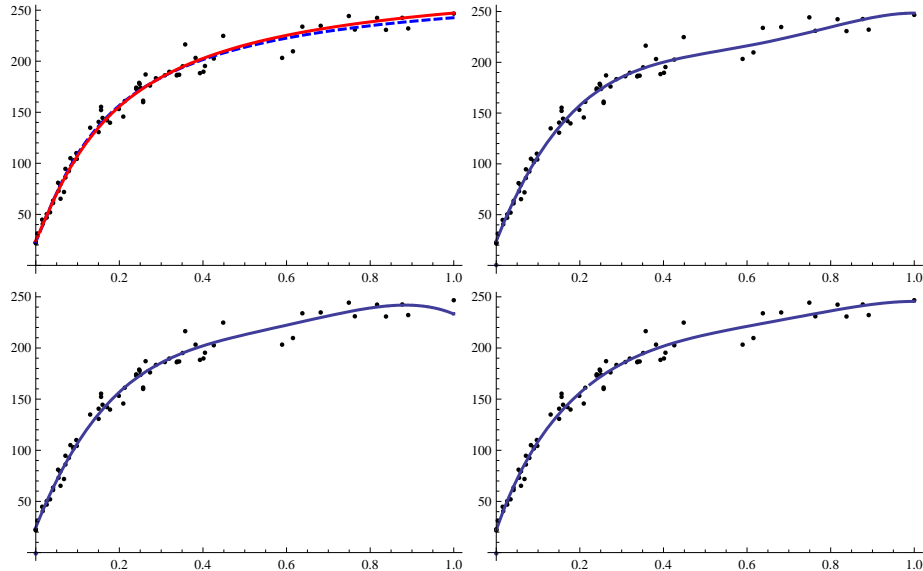


Figure 3.2: Parametric and shape-constrained spline estimators for the regression curves of the Rabbit data (scaled to $[0, 1]$). Top left: original estimator from [DM61] (red, solid), and the ordinary least squares estimator (blue, dashed). Top right: IQ estimator. Bottom left: VQ estimator. Bottom right: IVQ estimator. (See the beginning of the Section 3.5 for the explanation of acronyms.)

optimal. We compared these models to different shape-constrained polynomial spline models by measuring the L_2 , L_1 , and L_∞ distances of the true function f and the estimators, together with their estimated numbers of parameters and their AIC_c . This process was repeated 100 times. The means and standard deviations of the distances are reported in Table 3.3. In the table headings μ_p and σ_p denote the estimated mean and standard deviation of the L_p distances; $p \in \{1, 2, \infty\}$. Similarly, μ_{AIC} and σ_{AIC} are the mean and standard deviation of the AIC_c scores of the estimators. PM1 and PM2 are the least-squares solution to the parametric models above, the remaining acronyms in this table are explained at the beginning of Section 3.5.

model	$10^3 \cdot \mu_2$	$10^3 \cdot \sigma_2$	$10^3 \cdot \mu_1$	$10^3 \cdot \sigma_1$	μ_∞	σ_∞	μ_{AIC}	σ_{AIC}
PM1	0.972	1.516	15.685	12.268	0.054	0.050	-20.727	11.443
PM2	7.552	3.051	65.173	15.911	0.265	0.100	-12.934	11.513
NC	5.635	4.675	55.616	21.907	0.201	0.148	-13.523	11.418
NIC	2.768	2.992	36.890	18.938	0.133	0.099	-11.022	11.486
NVC	3.277	2.942	43.397	18.716	0.128	0.067	-11.826	11.423
NIVC	2.243	2.309	34.466	17.964	0.111	0.076	-10.602	11.480
NIVQ	2.613	2.567	36.431	17.667	0.138	0.102	-8.220	11.385

Table 3.3: Comparison of parametric and spline estimators for the regression curves of a simulated dataset. μ_i and σ_i are the estimated mean and standard deviation of the L_i error based on 100 experiments. The remaining acronyms in the table headings are explained in the beginning on Section 3.5.

It is immediate from Table 3.3 that the estimators benefit from each added shape constraint. Both the increasing and the concave cubic splines are considerably better estimates than the (otherwise

unconstrained) nonnegative splines, and the concave increasing spline is better than both of them, with respect to every distance measure considered. Obviously, the AIC_c scores worsen as we add more constraints to the model. Quartic splines, on the other hand, are not any better than cubic ones in this example. Notice that the best models outperform even the parametric model PM2 with respect to every measure of goodness-of-fit, except for AIC_c , the latter owing to the larger number of parameters. The optimal NIVC model usually had 5 or 6 effective parameters compared to the 2 fixed parameters of PM2. As expected, the single-parameter model PM1 proved to be hard to match.

3.5.2 Density estimation

The following tests were conducted to compare our methods to a wide range of kernel methods. The choice of benchmark distributions, as well as the experimental design is based on [EL01, Chapter 8].

The probability density functions of the benchmark distributions are:

$$f_1(x) = \frac{9}{10}\phi_{1/2}(x-5) + \frac{1}{10}\phi_{1/2}(x-7), \quad (3.14a)$$

$$f_2(x) = \phi_1(x-5), \quad (3.14b)$$

$$f_3(x) = \frac{1}{5}U([3, 8]), \quad (3.14c)$$

$$f_4(x) = \frac{1}{5}\psi_{1.4, 2.6}\left(\frac{1}{5}(x-0.3)\right) \quad (3.14d)$$

$$f_5(x) = \frac{1}{4}\phi_{9/5}(x-6) + \frac{4}{5}\phi_{1/10}(x-2) \quad (3.14e)$$

$$f_6(x) = \frac{1}{2}\phi_{1/2}(x-3.5) + \frac{1}{2}\phi_{1/2}(x-6.5), \quad (3.14f)$$

where $\phi_\sigma(x)$ is the pdf of the normal distribution with mean zero and standard deviation σ , $U([a, b])$ is the pdf of the uniform distribution on $[a, b]$, and $\psi_{\alpha, \beta}$ is the pdf of the Beta density with parameters α and β . (See Figure 3.3).

For completeness we included the uniform distribution from [EL01], even though f_3 is not continuous, hence the theoretical justification of our sieve methods (Section 3.2) does not apply. Another rather comprehensive, but less recent, simulation study of kernel density estimates can be found in [BD94].

Following [DG85], the authors of [EL01] argue that the most meaningful comparison between estimates of probability densities is the L_1 distance of their PDFs, and they report only this measure of goodness-of-fit in their benchmark. Hence in this section we do the same when comparing our methods and theirs.

For each benchmark density, random samples of size 100 were generated. Then the optimal cubic




Figure 3.3: Benchmark distributions (3.14) from [EL01].

spline densities were determined using the Bernstein polynomial based and the SOCP based methods outlined in Sections 3.3.1 and 3.3.2, and compared to the kernel estimates of [EL01], which employ the Epanechnikov kernel and the normal density kernel, and thirteen(!) bandwidth selection methods. These bandwidth selection methods include the “optimal method”, which simply determines the bandwidth that minimizes the L_1 error. This is clearly not a rational method, as it requires the knowledge of the estimated PDF, but it serves as a perfect benchmark, as it is an upper bound on the performance of all possible bandwidth selection methods. The process was repeated 100 times, the means and standard deviations of the L_1 distances of the estimators and the true PDFs are reported in Table 3.4.

To avoid repeating all the results of the book, we refer the reader to the tabulated numerical results of [EL01, pp. 326–327]. In our table we only show the L_1 errors of the “optimal method”. Columns OP-E and OP-N show the results of the optimal method using the Epanechnikov and the normal density kernel, respectively. The columns Bernstein and SOCP contain the results of the least-squares best fitting piecewise Bernstein polynomial and nonnegative cubic splines, respectively. The rows $\mu^{(i)}$ and $\sigma^{(i)}$ show the mean and the standard deviation of the L_1 errors from the experiments with the density f_i .

It is safe to conclude that in all but one of these examples both the Bernstein polynomial-based and the SOCP methods give results comparable to those of the kernel methods. Except for the density f_5 (a bad mixture of normal distributions) both methods nearly match, break even with, or outperform even the optimal kernel method. The results are easy to interpret: the less smooth the estimated PDF is, the higher the disadvantage of cubic splines to the kernel methods. The difference between Bernstein polynomial-based and the SOCP methods is too small to be considered significant, their performance is essentially identical in all the examples.

	OP-E	OP-N	Bernstein	SOCP
$\mu^{(1)}$	0.158	0.160	0.150	0.159
$\sigma^{(1)}$	0.053	0.053	0.049	0.045
$\mu^{(2)}$	0.127	0.129	0.064	0.075
$\sigma^{(2)}$	0.054	0.054	0.069	0.060
$\mu^{(3)}$	0.217	0.218	0.171	0.179
$\sigma^{(3)}$	0.042	0.042	0.069	0.075
$\mu^{(4)}$	0.162	0.164	0.148	0.153
$\sigma^{(4)}$	0.041	0.041	0.065	0.064
$\mu^{(5)}$	0.335	0.335	0.649	0.643
$\sigma^{(5)}$	0.055	0.055	0.022	0.051
$\mu^{(6)}$	0.187	0.189	0.227	0.226
$\sigma^{(6)}$	0.053	0.053	0.031	0.026

Table 3.4: Estimated means and standard deviations of the L_1 errors of kernel and spline estimates of the densities (3.14). $\mu^{(i)}$ and $\sigma^{(i)}$ are the mean and standard deviation of the L_1 errors from 100 experiments with f_i ; see the text for details and the interpretation of the column headings.

3.5.3 Arrival rate estimation

The paper [AENR08] gives detailed numerical evaluation of the SDP/SOCP approach and they show that it compares very favorably to some wavelet-based methods. We only compare the SDP/SOCP approach and the nonnegative basis approach using Bernstein polynomials; first using the same email arrival data that was used introduced in [AENR08], and then using simulated arrival data.

Email Arrival Data. We first present results obtained for a dataset of 10,150 e-mail arrivals recorded in a period of 446 days. As described in [AENR08], the arrival times were rounded to the nearest second. It was assumed that the arrival rate is periodic with a period of one week. The details of the appropriate maximum likelihood model using nonnegative cubic splines is also given in [AENR08], and using cross-validation the authors of [AENR08] found that the optimal number of evenly distributed knot points to be used is 48.

We determined the optimal nonnegative spline as well as the optimal piecewise Bernstein polynomial spline with $7k$, ($k = 3, \dots, 14$) knot points, and with the quoted 48 knot points. We found that with 28 knot points there was some small but visible difference between the two corresponding splines. Not surprisingly, the differences occur during night times (when the arrival rate is close to zero) and in the subsequent mornings. In all the other experiments, when the number of knot points were different from 28, the two splines were identical.

Simulated Data. Simulated experiments were carried out in the following manner. Arrival times were simulated by our implementation of the thinning method [LS78], using a simple arrival rate

function,

$$\lambda(t) = 4000(\cos(6\pi t + 0.2) + 2.16 \sin(4\pi t + 3.8) + 2.6),$$

which demonstrates several properties that can potentially make the approximation difficult: it gets near zero, it has a high and narrow peak, and it also has a long flat part. It was also chosen to be periodic with period 1, as most practical applications assume periodic arrival rate function.

Assuming that the interval $[0, 1]$ represents a week, arrival times were rounded to the nearest minute. Then the optimal cubic spline arrival rates were determined using the Bernstein polynomial based and the SOCP based methods outlined in Sections 3.3.1 and 3.3.2, with AIC_c used in knot point selection. This process was repeated a hundred times. Figure 3.4 shows the 95% empirical confidence bands plotted against each other and the true arrival rate function. The red curve is $\lambda(t)$, the blue band corresponds to the SOCP approach, and the green band corresponds to the nonnegative basis approach. It is apparent that the difference between the results of the two approaches is quite small. Figure 3.5 shows the widths of the confidence bands plotted against each other. It is clear from this plot that the quality of the SOCP-based model is nearly pointwise better than that of the nonnegative basis approach.




Figure 3.4: Comparison of the 95% empirical confidence bands of the results from SOCP and nonnegative basis approaches. The blue curve corresponds to the SOCP approach, and the green one to the piecewise Bernstein polynomial splines.

3.6 Summary and conclusion

Two optimization-based approaches using conic sieves of low-degree splines were presented, and compared to each other and to popular kernel methods in various applications of shape constrained estimation problems, including those with a mixture of shape constraints.

Theoretically, the SDP/SOCP approach, which requires solving optimization problems with linear and second-order cone constraints, is clearly superior to the basic nonnegative basis approach (which




Figure 3.5: The bandwidths of the 95% empirical confidence bands of the results from SOCP and nonnegative basis approaches. The blue curve corresponds to the SOCP approach, and the green one to the piecewise Bernstein polynomial splines.

employs only linear constraints to represent polyhedral sieves), as in univariate problems it allows us to optimize precisely over the set that we want: the set of nonnegative splines. Using polyhedral sieves results in simpler, linearly constrained models, but in most cases it only allows us to optimize over a strictly smaller or larger set than is necessary.

In many instances the two approaches gave very similar results, but we have also found several examples where the SOCP approach is superior in practice, too. On the other hand, we exhibited a few cases when the two approaches are provably equivalent.

The optimization models obtained from both approaches were easily solvable in a fraction of a second using readily available nonlinear solvers. From the viewpoint of computational feasibility, these approaches are competitive even with closed-form formulae, which are only available for very restricted special cases.

Previously it has been established that for arrival rate estimation our approach yields superior results to methods using wavelets. In our experiment we found that we can also match, and frequently outperform state-of-the-art kernel methods in density estimation problems.

Chapter 4

Multivariate function estimation with shape constraints

Our aim in this chapter is to extend the ideas introduced in the previous chapter to multivariate estimation problems. This extension is far from straightforward for a number of reasons.

- First, the nonnegativity constraint becomes formally “intractable” for multivariate problems, and is very difficult to handle even in problems involving bi- or trivariate splines of low degree.
- Some of the most basic constraints on the shape of the estimator will no longer reduce conveniently to nonnegativity. Convexity is perhaps the simplest example, which directly translates (for twice differentiable functions) to the positive semidefiniteness of the Hessian. Hence, we need a more general notion of nonnegativity that brings shape constraints of multivariate functions under one umbrella. We investigate this problem in Section 4.3. As an additional benefit, the tools we develop for such problems also allows us to revisit some *univariate* problems we were not able to solve before.
- Finally, we shall note that our methods (or any other spline-based method for that matter) cannot be applied to very high-dimensional problems, as the number of pieces the estimator consists of inevitably becomes enormous. In fact, solving the proposed optimization problems becomes difficult, if not impossible, even for some problems of very low dimension, owing to the size of the problem. In Section 4.2 we shall explore the application of decomposition methods that exploit the structure of our models and enable us to solve very large-scale spline estimation problems.

4.1 Multivariate nonnegative and sum-of-squares splines

In this section we outline how conic sieves of shape-constrained splines can be applied to solve multivariate shape-constrained statistical estimation problems when the shape constraints reduce to nonnegativity. This topic has been studied less than the univariate estimation problems of Chapter 3, and the existing methods are typically direct generalizations of the heuristic methods common in the univariate literature. For example, when investigating a bivariate estimation problem with bound constraints on the estimator, Villalobos and Wahba [VW87] suggest that the constraints $0 \leq s(\mathbf{x}) \leq 1$ ($\forall \mathbf{x} \in \Delta$) on the spline estimator s be replaced by $0 \leq s(\mathbf{x}_i) \leq 1$ ($i = 1, \dots, N$), where \mathbf{x}_i are the points of a fine enough grid in Δ . However, there is no guarantee that the resulting estimator will indeed satisfy the bounds everywhere regardless of how the grid points \mathbf{x}_i are chosen. At best, the right mesh-size must be found by trial and error.

If instead we wish to tackle the nonnegativity constraint directly, we must remember that deciding whether a multivariate polynomial is nonnegative over a given polyhedral region is NP-hard, even if the region is a very simple fixed region, such as the unit cube, even for degree two (Proposition 2.11). If both the number of variables and the degree are kept small (and are considered constant), the recognition of nonnegative polynomials becomes polynomial-time solvable [Ren92], but it requires further investigation how the resulting quantifier-free first-order representations of nonnegative polynomials can be incorporated into a tractable optimization model. (They generally translate to a large collection of non-convex constraints.) However, Section 3.3 provides us with two approaches that can be generalized to the multivariate setting:

1. Following Section 3.3.2, we may consider a basis of polynomials or splines which consists of nonnegative functions, and confine the search for an optimal estimator to the nonnegative linear combination of these. Specifically, we may use multivariate Bernstein polynomials.
2. Following Proposition 2.5, it has been proposed in the constrained polynomial programming literature that constraints of the form

$$p(\mathbf{x}) \geq 0 \quad \forall \mathbf{x} \in X = \{\mathbf{x} \mid g_i(\mathbf{x}) \geq 0, i = 1, \dots, I\},$$

where g_1, \dots, g_I are given polynomials (and p is the unknown polynomial) be replaced by the (stronger) constraints

$$p = \sum_{i=1}^I g_i \sum_{j=1}^{N_i} r_{ij}^2 \quad (r_{ij} \in V_i), \quad (4.1)$$

where V_1, \dots, V_I are some fixed (finite-dimensional) linear spaces of polynomials. Polynomials p of the form (4.1) are called *weighted-sum-of-squares* (or *WSOS*) polynomials; recall also the theorems of Schmüdgen and Putinar, and the representations (2.4) and (2.5).

It is perhaps worth mentioning that the first approach is the special case of the second approach, when $\{g_i\}$ is chosen to be a nonnegative basis of polynomials, and each V_i is the one-dimensional space of constant polynomials.

We may also consider splines that are piecewise WSOS polynomials. The nonnegativity of a spline s can be imposed by considering piecewise WSOS polynomial splines, where the weights g_i of each piece are chosen to be (usually affine) polynomials that define the domain of the piece.

Regardless of how we choose the weights and the spaces V_i , constraints of the form (4.1) on p can be translated to linear and semidefinite programming (SDP) constraints on the coefficients of p [Nes00], just like in the univariate case. However, in the multivariate setting there is no analogue of Proposition 2.5: the cone of SOS multivariate polynomials is always a *proper* subset of nonnegative polynomials, aside from the cases mentioned in Proposition 2.10. Similarly, cones of piecewise WSOS splines are proper subsets of nonnegative splines. Since we have complete freedom in choosing the subdivision of the domain as well as the spaces V_i in defining the weighted-sum-of-squares cone in (4.1), care needs to be taken so the piecewise WSOS polynomial splines used in our conic sieve are dense in some “rich enough” space of functions \mathbb{H} , to which the ideal estimator s is assumed, or known, to belong. In the remaining of this section we give a practical sufficient condition when this is the case.

To keep things simple, let us assume that Δ is the m -dimensional interval $[0, 1]^m$, and consider splines of total degree d over a rectilinear subdivision of Δ . Formally, a rectilinear subdivision of Δ is given by a list of vectors (of possibly different dimensions) $\mathbf{A} = (a_{i,j}), i = 1, \dots, d, j = 1, \dots, \ell_i$ such that the knot point sequence $(a_{i,1}, \dots, a_{i,\ell_i})$ defines the subdivision of Δ along the i th axis. Then each region of the subdivision is affinely similar to Δ , and we can represent a spline by the coefficients of its polynomial pieces scaled to Δ .

For each dimension $i = 1, \dots, d$ we fix a basis $\{u_0^{(i)}, \dots, u_{m_i}^{(i)}\}$ of polynomials of degree m_i , with domain $[0, 1]$. The spline s is then given piecewise by

$$s(\mathbf{x}) = q^{(\mathbf{j})}(\mathbf{x}) \quad \forall \mathbf{x} = [a_{1,j_1}, a_{1,j_1+1}] \times \dots \times [a_{d,j_d}, a_{d,j_d+1}],$$

for each multi-index $\mathbf{j} = (j_1, \dots, j_d)$; and each polynomial piece $q^{(\mathbf{j})}$ is represented by the coefficients

$p_{\mathbf{k}}^{(\mathbf{j})}$, $\mathbf{k} \in \{0, \dots, m_1\} \times \dots \times \{0, \dots, m_k\}$, of its affinely scaled counterpart $p^{(\mathbf{j})} : [0, 1]^d \rightarrow \mathbb{R}$ satisfying

$$q^{(\mathbf{j})}(\mathbf{x}) = \sum_{\mathbf{k}} p_{\mathbf{k}}^{(\mathbf{j})} \prod_{i=1}^d u_{k_i}^{(i)} \left(\frac{x_i - a_{i,j_i}}{a_{i,j_i+1} - a_{i,j_i}} \right). \quad (4.2)$$

It is clear that $s(\mathbf{x}) \geq 0$ for every $\mathbf{x} \in \Delta$ if and only if $p^{(\mathbf{j})}(\mathbf{x}) \geq 0$ for every \mathbf{j} and $\mathbf{x} \in \Delta$. We refer to this representation of s by the coefficients $p_{\mathbf{k}}^{(\mathbf{j})}$ as the *scaled representation* of s .

The *mesh size* of a subdivision \mathbf{A} is defined as $\|\mathbf{A}\| = \max_{i,j}(a_{i,j+1} - a_{i,j})$. Nestedness of subdivisions can also be defined analogously to the univariate case. Let us fix g_i and V_i , and let us denote by W the convex cone of piecewise WSOS polynomials over Δ defined by (4.1). For a given subdivision \mathbf{A} and WSOS set W , $\mathcal{P}(W, \mathbf{A})$ denotes the set of piecewise WSOS polynomial splines over the subdivision \mathbf{A} whose pieces (when scaled to Δ) belong to W . Again, $\mathcal{P}(W, \mathbf{A})$ is a convex cone. We have the following theorem.

Theorem 4.1. *Assume that $1 \in \text{int } W$, where 1 denotes the constant one polynomial. Furthermore, let $\{\mathbf{A}_i\}$ be an asymptotically nested sequence of subdivisions of Δ satisfying $\lim \|\mathbf{A}_i\| = 0$. Then the set $\bigcup_i \mathcal{P}(W, \mathbf{A}_i)$ is a dense subcone of $\mathcal{P} \cap C(\Delta)$, the cone of nonnegative continuous functions over Δ .*

We need not prove this theorem directly; instead, we can prove, analogously to Theorem 3.4, a stronger assertion using polyhedral cones; see Theorem 4.2 below.

A sufficient condition for a polynomial to be nonnegative over Δ is for it to have nonnegative coefficients in a basis $U = \{u_0, \dots, u_m\}$ that consists of polynomials nonnegative over Δ , that is, for it to belong to $\text{cone}(U)$ for a nonnegative basis U . Similarly to the piecewise WSOS polynomial splines above we can define a *piecewise U -spline* as a piecewise polynomial spline whose pieces (in the scaled representation) belong to $\text{cone}(U)$. The set of piecewise U -splines with subdivision \mathbf{A} is denoted by $\mathcal{P}(U, \mathbf{A})$.

Theorem 4.2. *Consider a basis $U = \{u_0, \dots, u_m\}$ of d -variate polynomials of multi-degree $\mathbf{m} = (m_1, \dots, m_d)$ such that each u_i is nonnegative over $\Delta = [0, 1]^d$, and assume that $1 \in \text{int } \text{cone}(U)$, where 1 denotes the constant one polynomial. Furthermore, let $\mathbf{A}_1, \mathbf{A}_2, \dots$ be an asymptotically nested sequence of subdivisions with mesh sizes approaching zero. Then the set $\bigcup_i \mathcal{P}(U, \mathbf{A}_i)$ is a dense subcone of the cone of nonnegative functions over Δ .*

Proof. First we show that for every polynomial p of degree \mathbf{m} , strictly positive over $[0, 1]$, there exist nonnegative constants C_i such that $p + C_i \in \mathcal{P}(U, \mathbf{A}_i)$ for every i , and $\lim C_i = 0$.

Fix i , and consider a piece in the subdivision from the knot point sequence \mathbf{A}_i :

$$[a_{1,j_1}, a_{1,j_1+1}] \times \cdots \times [a_{d,j_d}, a_{d,j_d+1}].$$

The polynomial p can be represented as a piecewise polynomial spline of degree \mathbf{m} with knot point sequence \mathbf{A}_i ; its scaled representation is

$$p^{(\mathbf{j})}(x_1, \dots, x_d) = p((a_{1,j_1+1} - a_{1,j_1})x_1 + a_{1,j_1}, \dots, (a_{d,j_d+1} - a_{d,j_d})x_d + a_{d,j_d}).$$

Collecting terms in the standard basis, we see that every coefficient in the above expression is of order $\mathcal{O}(\|\mathbf{A}_i\|)$, except for the constant term, which is $p(a_{1,j_1}, \dots, a_{d,j_d})$. By assumption, this constant term is positive, because p is strictly positive on $[0, 1]$. By the assumption on U , $\sum_{k=0}^m \alpha_k u_k \equiv p(a_{1,j_1}, \dots, a_{d,j_d})$ for some positive $\alpha_0, \dots, \alpha_m$. Now, if we express $p^{(\mathbf{j})}$ in the basis U : $p^{(\mathbf{j})} = \sum p_k^{(\mathbf{j})} u_k$, we have that $p_k^{(\mathbf{j})} = \alpha_k - \delta_k^{(\mathbf{j})}$ with $|\delta_k^{(\mathbf{j})}| = \mathcal{O}(\|\mathbf{A}_i\|)$, consequently $p^{(\mathbf{j})} + p(a_{1,j_1}, \dots, a_{d,j_d}) \max_k (|\delta_k^{(\mathbf{j})}| / \alpha_k)$ has positive coefficients in the basis U . Applying the same argument for every \mathbf{j} , we obtain that $p + C_i \in \mathcal{P}(U, \mathbf{A}_i)$ for

$$C_i = \max_{\mathbf{j}} (p(a_{1,j_1}, \dots, a_{d,j_d}) \max_k (|\delta_k^{(\mathbf{j})}| / \alpha_k)).$$

Finally, as $|\delta_k^{(\mathbf{j})}| = \mathcal{O}(\|\mathbf{A}_i\|)$ and p is bounded, $C_i \rightarrow 0$ as $\|\mathbf{A}_i\| \rightarrow 0$.

The same argument also proves that for every strictly positive spline over $[0, 1]$, with knot point sequence \mathbf{A} , and for every sequence $\{\mathbf{A}_i\}$ consisting of subdivisions of \mathbf{A} satisfying $\lim \|\mathbf{A}_i\| = 0$, there exist nonnegative constants C_i such that $s + C_i \in \mathcal{P}(U, \mathbf{A}_i)$ for every i , and $\lim C_i = 0$.

Consequently, $\bigcup_i \mathcal{P}(U, \mathbf{A}_i)$ is a dense subset of nonnegative splines of multi-degree \mathbf{m} .

Now our assertion follows from the fact that tensor product splines (of every given order of differentiability) in $\Delta = [0, 1]^d$ are dense in the space of continuous functions over Δ ; see [Sch81, Theorem 13.21]. \square

Note that the conditions $1 \in \text{int } \Sigma$ and $1 \in \text{int cone}(U)$ are sufficient *and necessary* for the desired conclusion. For example, the cone of polynomials with nonnegative coefficients in the standard monomial basis is a WSOS cone with weights nonnegative over $[0, 1]^d$. It does not satisfy the condition, as the constant 1 is on the boundary of this polynomial cone. Incidentally, the corresponding cone of splines consists of functions that are monotone nondecreasing in every variable, hence it is not dense in the cone of nonnegative continuous functions over $[0, 1]^d$.

As in the univariate case, a simple example is provided by the multivariate Bernstein polynomials.

(Multivariate Bernstein polynomials are of the form $\prod_i B_{n_i, k_i}(x_i)$, where each B_{n_i, k_i} is a univariate Bernstein polynomial.) The cone of polynomials whose coefficients are nonnegative in the Bernstein polynomial basis contains 1 in its interior, hence the piecewise Bernstein polynomial splines are dense in the cone of nonnegative functions in $C(\Delta)$.

As a final remark, we shall clarify in what sense the second, polyhedral approximation, approach is a special case of the WSOS approach. The cone $\text{cone}(U)$ for a nonnegative polynomial basis U can be considered a WSOS cone defined in (4.1) with weights in U , whose spaces V_i are the one-dimensional linear spaces consisting only of constant polynomials. Furthermore, for every WSOS cone satisfying the conditions of Theorem 4.1 one can find a basis U such that the corresponding polyhedral spline cone satisfies the conditions of Theorem 4.2.

4.1.1 Numerical illustration – posterior probability estimation

We conclude this section with a numerical illustration. This problem is borrowed from Villalobos and Wahba [VW87], and is an extension of two-class classification: posterior probability estimation.

Consider two m -dimensional random variables: X_1 and X_2 with (unknown) density functions f_1 and f_2 , respectively. The random vector variable Y is an observation from either X_1 or X_2 ; if it is from X_i , it has density function f_i . Let p_i be the (known or unknown) prior probability that Y belongs to X_i . Our goal is to estimate the *posterior probability* of X_1 ,

$$f(\mathbf{y}) \stackrel{\text{def}}{=} \Pr(X_1 | Y = \mathbf{y}) = \frac{p_1 f_1(\mathbf{y})}{p_1 f_1(\mathbf{y}) + p_2 f_2(\mathbf{y})},$$

based on a collection of i.i.d. samples from Y , $\{\mathbf{y}_1, \dots, \mathbf{y}_n\}$ with known class information: $z_i = 0$ if and only if \mathbf{y}_i belongs to X_1 , otherwise $z_i = 1$.

If the prior probabilities p_i are unknown (which is the case in most applications), we can estimate them by the proportion of samples that belong to population X_i . Denoting these proportions by \hat{p}_i , we consider the problem of estimating

$$\hat{f}(\mathbf{y}) = \frac{\hat{p}_1 f_1(\mathbf{y})}{\hat{p}_1 f_1(\mathbf{y}) + \hat{p}_2 f_2(\mathbf{y})}.$$

Following [VW87], observe that if Z is the indicator variable that assumes value $i - 1$ when Y is from population i , then $\mathbb{E}(Z|Y = \mathbf{y}) = \hat{f}(\mathbf{y})$. Hence, we can think of values z_i as “noisy” values of

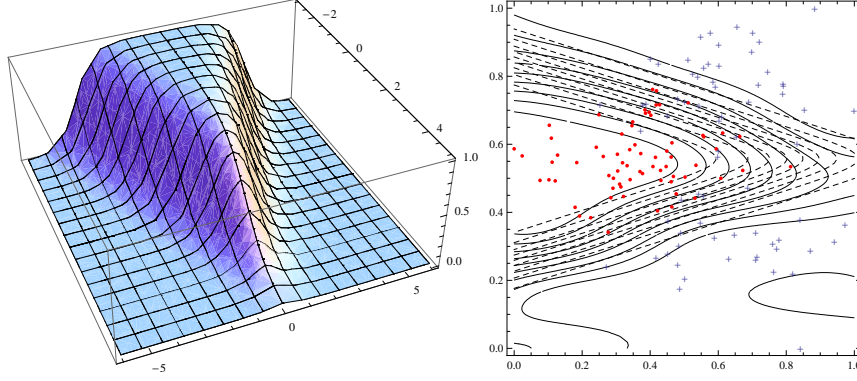


Figure 4.1: Posterior probability estimation example. Left: the true posterior probability. Right: the contour plots of the true posterior probability (dashed) and the smoothing spline estimator (solid). The blue dots and red ‘+’ signs are samples from the two populations (scaled to $[0, 1]^2$).

$\hat{f}(\mathbf{y}_i)$. This motivates the following optimization models: find a spline $s: \Delta \rightarrow \mathbb{R}$ that satisfies

$$0 \leq s(\mathbf{y}) \leq 1 \quad \forall \mathbf{y} \in \Delta, \quad (4.3)$$

and which minimizes the residual sum of squares loss function $L(s) = \sum_{i=1}^n (z_i - s(\mathbf{y}_i))^2$, or the loss function with smoothing penalty,

$$L_\lambda(s) = \sum_{i=1}^n (z_i - s(\mathbf{y}_i))^2 + \lambda J_k(s) \quad (\lambda > 0),$$

where

$$J_k(s) = \sum_{\alpha_1 + \dots + \alpha_m = k} \frac{k!}{\alpha_1! \dots \alpha_m!} \int \dots \int \left(\frac{\partial^k s}{\partial y_1^{\alpha_1} \dots \partial y_m^{\alpha_m}} \right)^2 dy_1 \dots dy_m.$$

It is easy to show that J_k is a convex quadratic function of s if s is a polynomial of total degree k . In particular, minimizing $L_\lambda(s)$ where s is a multivariate spline over a fixed subdivision of the domain Δ amounts to solving a convex quadratic optimization problem. By applying piecewise weighted-sum-of-squares approximation to (4.3), we obtain a semidefinite-representable problem.

Turning to the numerical example of [VW87], let $N(\mu_1, \mu_2; \sigma_1, \sigma_2)$ denote the bivariate uncorrelated normal distribution with mean (μ_1, μ_2) and covariance matrix $\text{diag}(\sigma_1, \sigma_2)$, and let $p_1 = p_2 = 1/2$, and

$$f_1 \sim N(0, 0; 1, 1) \quad \text{and} \quad f_2 \sim \frac{1}{2}N(1.5, -2.5; 1, 1) + \frac{1}{2}N(1.5, 2.5; 1, 1).$$

The true posterior probability function is shown on the left-hand side of Figure 4.1.

We generated 70 samples from both populations, and scaled them so that the axis-parallel bounding box of the samples became the unit square $[0, 1]^2$. The right-hand side of Figure 4.1 shows the sample points and the contours of the true function and of the optimal piecewise Bernstein polynomial smoothing spline, with uniform subdivision of $[0, 1]^2$. The optimal number of pieces, as well as the smoothing parameter λ were determined by cross-validation.

In this example, the estimator follows the contours of the true function, but it is smoother than the true function. The example also shows a drawback of the using a uniform rectangular subdivision of the domain: unexpected behavior of the estimator may occur in regions where there are no samples. This is a problem also to be expected in regression problems (where the extrapolation is meaningless), but can be somewhat remedied by smoothing techniques, or by bounding the derivatives of the estimator at the boundary, as in *natural splines* [Wah90], or by using subdivisions of the domain that are adapted to the data. The effects of choosing different subdivisions (which also necessitates a generalization of Theorem 4.1) is subject to future research.

4.2 A decomposition method for multivariate splines

As we mentioned in the introduction of this chapter, the size of the optimization models involving multivariate splines prohibits the solution of these models when the mesh size is small. On the other hand, these problems have a very regular and sparse structure amenable to decomposition methods. In this section we outline an augmented Lagrangian decomposition method with particularly good convergence properties for spline estimation problems, and illustrate it by estimating the (two-dimensional) weekly arrival rate of car accidents on the New Jersey Turnpike.

There is a vast literature on decomposition methods for both linear and nonlinear convex optimization problems, an area initiated by Dantzig and Wolfe [DW60] and Benders [Ben62], and we can only choose to give a very brief review of the fundamental ideas relevant in our context.

Most methods are directly based on Lagrangian duality (see, for example, [Las70, Chapter 8] or [Ber99, Chapter 4] for an overview of the basic methods). They associate Lagrange multipliers to the linear “complicating” or “coupling” constraints, and solve the resulting Lagrangian relaxation, which is automatically decomposable. The different methods primarily differ in their strategies to iteratively update the Lagrange multipliers until the optimal multipliers, with which the Lagrangian relaxation is equivalent to the original problem, is found. Drawbacks of these methods include the non-differentiability of the dual functional and the problem that reconstruction of the solutions of the original problem from the optimal solution of the Lagrangian relaxation is not trivial [Rus95].

Replacing the Lagrangian dual by the *augmented Lagrangian*, as in [Hes69], which includes in the objective an extra quadratic penalty term for constraint violation, solves these problems, and has the added advantage of making the dual multiplier updates very easy [Rus05]. The same reasons motivate *proximal point methods*, which also use a quadratic regularizing term. However, the quadratic terms introduced by these approaches are not decomposable even if the constraints in the problem are linear. Hence, several methods have emerged based either on linearizations of the augmented Lagrangian, such as [WNM76, Rus95], or on linearization in proximal point methods [Roc76, KRR99].

Next, we shall describe the method we shall use for spline estimation problems. It is a simplified version of the augmented Lagrangian-based method from [Rus95] specifically designed for sparse problems. Since some of the details of our specific problems might obscure the main ideas of the algorithm, we shall discuss the method in a slightly more abstract form than necessary for our purposes.

4.2.1 Augmented Lagrangian decomposition for sparse problems

Let $L \geq 2$, and let \mathcal{X}_i ($i = 1, \dots, L$) be a nonempty compact subset of \mathbb{R}^{n_i} . Finally, let $f_i: \mathcal{X}_i \rightarrow \mathbb{R}$ be convex. With these given, we consider the convex optimization problem

$$\text{minimize } f(\mathbf{x}) \stackrel{\text{def}}{=} \sum_{i=1}^L f_i(\mathbf{x}_i) \tag{4.4a}$$

$$\text{subject to } \sum_{i=1}^L \mathbf{A}_i \mathbf{x}_i = \mathbf{b} \tag{4.4b}$$

$$\mathbf{x}_i \in \mathcal{X}_i \quad i = 1, \dots, L. \tag{4.4c}$$

As a specific example, we can model shape-constrained optimization problems in this framework: the optimality of the estimator is typically defined piecewise, and so are the shape constraints, under some continuity assumptions. (For example, if the estimator is assumed to be continuous, then it is monotone increasing if and only if it is piecewise monotone increasing.) Thus, we can set L to be the number of pieces, (4.4c) are the constraints on the shape of the estimator (now defined piecewise), and (4.4b) includes the continuity of the estimator and perhaps the continuity of its derivatives, as well as periodicity, interpolation constraints, as needed for the problem.

Alternatively, if the estimator is a polynomial spline over a rectilinear grid we can set $L = 2$, since all constraints connect only pieces of two disjoint class, following a chessboard-like pattern. The same is true for some other regular subdivisions as well, including a regular simplicial subdivision.

This property can be exploited by some methods. We shall focus on a direct consequence of this observation: that in our spline estimation problems *all of the coupling constraints in (4.4b) involve variables corresponding to only two different \mathcal{X}_i* . We must immediately add the only exception to the rule: density estimation problems (Section 3.4.3) include the constraint that the estimator integrates to one, which is a single linear equation that involves every variable. (Arrival rate estimation problems also involve an integral, but in the objective function, which decomposes as a sum of integrals, computed piecewise.)

The condition that the sets \mathcal{X}_i be bounded is a rather technical condition, as one can always find reasonable bounds on the spline coefficients based on the data. In problems where the integral of the estimator is explicitly given, this is easy, and it is also straightforward in regression problems when the estimator is fitted to a finite number of points. The case of arrival rate estimation is more interesting, we shall return to this question in Section 4.2.2.

The method proposed in [Rus95] associates multipliers $\boldsymbol{\pi}$ to the linear coupling constraints, and considers a separable approximation Λ_{apx} of the augmented Lagrangian of (4.4),

$$\Lambda(\mathbf{x}, \boldsymbol{\pi}) = f(\mathbf{x}) + \langle \boldsymbol{\pi}, \mathbf{b} - \mathbf{A}\mathbf{x} \rangle + \frac{\varrho}{2} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|^2,$$

in which the bilinear terms in the quadratic penalties are linearized around a point $\tilde{\mathbf{x}} \in \mathbb{R}^{\sum_{i=1}^L n_i}$:

$$\Lambda_{\text{apx}}(\mathbf{x}, \tilde{\mathbf{x}}, \boldsymbol{\pi}) \stackrel{\text{def}}{=} \sum_{i=1}^L \Lambda_i(\mathbf{x}_i, \tilde{\mathbf{x}}_i, \boldsymbol{\pi}) \stackrel{\text{def}}{=} \sum_{i=1}^L f_i(\mathbf{x}_i) - \langle \mathbf{A}_i^T \boldsymbol{\pi}, \mathbf{x}_i \rangle + \frac{\varrho}{2} \left\| \mathbf{b} - \mathbf{A}_i \mathbf{x}_i - \sum_{j \neq i} \mathbf{A}_j \tilde{\mathbf{x}}_j \right\|^2. \quad (4.5)$$

The approach then is to fix a set of multipliers, find an approximate minimizer of the corresponding augmented Lagrangian problem by iteratively minimizing Λ_{apx} and updating $\tilde{\mathbf{x}}$ so as to approximate the optimal solution better. The components Λ_i of the approximate Lagrangian can be optimized separately, and even in parallel (in a Jacobi, rather than a Gauss–Seidel, fashion), which is a very desirable property when infrastructure for massively parallel computations is available. (For the author’s experiments it was not.) Once an approximately optimal solution to the augmented Lagrangian is found, we can update the multipliers as in the original multiplier method. The formal definition of the method is given in Algorithm 1, which requires two parameters: the augmented Lagrangian coefficient ϱ , and a step size parameter τ .

It can be shown that for every $\varrho > 0$ and $0 < \tau < 1/(N - 1)$ Algorithm 1 is convergent. For the inner loop this is a special case of Theorem 1 in [Rus95], and for the outer loop this follows from the

Algorithm 1: Simplified augmented Lagrangian decomposition

```

parameters:  $\rho > 0, \tau > 0$ 
1 initialize  $\boldsymbol{\pi}$                                 /* arbitrary initial value */
2 repeat
3    $\boldsymbol{\pi} \leftarrow \boldsymbol{\pi} + \rho(\mathbf{b} - \mathbf{Ax})$ 
4   foreach  $i = 1, \dots, L$  do solve  $\min_{\mathbf{x}_i \in \mathcal{X}_i} \Lambda_i(\mathbf{x}_i, \tilde{\mathbf{x}}, \boldsymbol{\pi})$           /* parallel */
5   if  $\mathbf{A}_i \mathbf{x}_i \neq \mathbf{A}_i \tilde{\mathbf{x}}_i$  for any  $i = 1, \dots, L$  then
6      $\tilde{\mathbf{x}} \leftarrow \tilde{\mathbf{x}} + \tau(\mathbf{x} - \tilde{\mathbf{x}})$ 
7     go to step 4
8   end if
9 until  $\mathbf{Ax} = \mathbf{b}$ 
10 return  $\mathbf{x}$ 

```

convergence of the method of multipliers [Roc76]. Perhaps the most attractive feature of Algorithm 1 is that its speed of convergence depends very highly on the largest number N of variable blocks \mathbf{x}_i linked by a coupling equality constraint, favoring problems with a sparse neighborhood structure. (In our spline estimation examples $N = 2$, and generally we have $2 \leq N \leq L$.) Let us measure the progress of the algorithm by the difference

$$\Delta^{(k)} \stackrel{\text{def}}{=} \Lambda(\tilde{\mathbf{x}}^{(k)}, \boldsymbol{\pi}^{(k)}) - \min_{\mathbf{x} \in \mathcal{X}} \Lambda(\mathbf{x}, \boldsymbol{\pi}^{(k)}),$$

where the superscript k refers to the number of times the outer loop of Algorithm 1 has been executed. The following theorem establishes the rate of convergence under a technical assumption.

Theorem 4.3. *Assume that there is a $\gamma > 0$ such that for every $\boldsymbol{\pi}$ and every $\mathbf{x} \in \mathcal{X}$,*

$$\Lambda(\mathbf{x}, \boldsymbol{\pi}) - \min_{\mathbf{x} \in \mathcal{X}} \Lambda(\mathbf{x}, \boldsymbol{\pi}) \geq \gamma \text{dist}(\mathbf{x}, \arg \min_{\mathbf{x} \in \mathcal{X}} \Lambda(\mathbf{x}, \boldsymbol{\pi}))^2,$$

and let $\alpha = \max_i \|\mathbf{A}_i\|_2$. Then for every $\rho > 0$ and $0 < \tau < 1/(N - 1)$,

$$\Delta^{(k+1)} \leq \left(1 - \frac{\tau(1 - \tau(N - 1))}{1 + 2\rho\alpha^2(N - 1)^2\gamma^{-1}} \right) \Delta^{(k)}. \quad (4.6)$$

Proof. This is a special case of Theorem 2 in [Rus95]. □

The first condition of the theorem is the *quadratic growth condition* of [Rus95], and is satisfied by every problem considered in this thesis. Theorem 4.3 also suggests a way to set the parameter τ : minimizing the coefficient on the right-hand side of (4.6) we obtain that $\tau = (2(N - 1))^{-1}$ is the recommended choice, regardless of ρ . (The theory does not provide guidance in the selection of ρ .)

4.2.2 Numerical illustration – NJ Turnpike accidents

We consider the two-dimensional point process of car accidents on the New Jersey Turnpike (NJTP). The two dimensions are time (with an assumed weekly periodicity) and location along the road. It is not entirely clear whether this is indeed a Poisson process, as accidents may change the traffic pattern, which in turn affects the distribution of the accidents. Furthermore, coincidences in the location (such as accidents at a temporary construction site) have likelihood zero in every Poisson model. However, as the accidents are relatively rare (serious accidents that change the traffic pattern for a long period of time are even more so), and major highways are assumed to have no easy-to-hit objects, a Poisson model may be a reasonable approximation. The fact that accidents may occur more frequently close to exits does not contradict the non-homogeneous Poisson model.

The data. We obtained car accident data from the New Jersey Department of Transportation. The raw data contained information on every car accident in 2009 recorded at the accident locations by police officers. The time of the accident is rounded to the nearest minute, but it is not clear whether the recorded time is the approximate time of the accident, the time the police were notified of the accident, or the time the officers attended to the accident. Hence, we can consider this as noisy data, despite the apparent precision of the time data. The location is given by the Standard Route Identifier of the road segment and an approximate milepost reading (variably rounded, apparently to the nearest 0.05 mile or to the nearest mile).

We removed all entries from the data that corresponded to accidents in roads other than the NJTP segment marked I-95. This is an approximately 78 miles long segment stretching between two state borders (with Pennsylvania and New York, respectively) with no forks or joins. We also removed entries with missing milepost information. Date and time were present for every entry. While we could take these accidents into account directly in a maximum-likelihood approach (and their time and date information we shall not discard), such incomplete entries were few, and it is reasonable to assume that accidents whose milepost is not recorded follow the same milepost and time distribution as the entries with complete information. Hence, we simplify our model, and estimate the arrival rate based only on the entries with specified milepost. We can then divide the obtained arrival rate with the ratio of entries with complete information to account for accidents discarded because of missing milepost value. We also removed all accidents that happened on ramps while entering or leaving the highway, as they are confounding in multiple ways. Finally, this left us with 4138 accidents.

The optimization model. Setting up the optimization model is relatively easy. The nonnegativity constraints for the spline estimator are replaced by weighted-sum-of-squares constraints, exactly as in Section 4.1.1. As objective, we seek to maximize a log-likelihood function detailed below. (Since we have a large number of arrivals, using the maximum likelihood approach is justified.) Furthermore we assume that the arrival rate is periodic with a weekly period.

Since exact arrival information is not available (both coordinates are rounded), we consider the following setup, similar to the univariate aggregated model of [AENR08]: the possible range of arrivals Δ is divided into small regions (representing the arrivals that would be rounded to the same point), and the data is the number of arrivals within each region. Suppose the data are aggregated in regions Δ_i , $i \in \mathcal{I}$, and let the number of arrivals recorded in region i be n_i . Finally, let $N = \sum_{i \in \mathcal{I}} n_i$. Then in the non-homogeneous Poisson model the likelihood associated with the arrival rate λ is

$$\begin{aligned} L_{\mathbf{n}}(\lambda) &= \Pr(\text{number of arrivals} = N \mid \lambda) \cdot \Pr(\text{distribution of exactly } N \text{ arrivals} = \mathbf{n} \mid \lambda) \\ &= \frac{I^N}{N!} e^{-I} \cdot \frac{N!}{\prod_{i \in \mathcal{I}} n_i!} \prod_{i \in \mathcal{I}} \left(\frac{\int_{\Delta_i} \lambda}{I} \right)^{n_i} = \frac{e^{-I}}{\prod_{i \in \mathcal{I}} n_i!} \prod_{i \in \mathcal{I}} \left(\int_{\Delta_i} \lambda \right)^{n_i}, \quad \text{where } I = \int_{\Delta} \lambda, \end{aligned}$$

provided all the integrals in the above formulae exist (which they do, of course, if λ is a polynomial spline). Hence, maximizing the likelihood function is equivalent to maximizing

$$f(\lambda) \stackrel{\text{def}}{=} \ln \left(\left(\prod_{i \in \mathcal{I}} n_i! \right) L_{\mathbf{n}}(\lambda) \right) = - \int_{\Delta} \lambda + \sum_{i \in \mathcal{I}} n_i \ln \int_{\Delta_i} \lambda. \quad (4.7)$$

This is the objective function of our optimization model. As the following observation shows, further simplification is possible.

Lemma 4.4. *Let \mathcal{K} be a cone of nonnegative functions over Δ whose restrictions to each Δ_i are integrable. Define $f: \mathcal{K} \rightarrow \mathbb{R}$ as in (4.7), and assume that there exists a $\lambda_0 \in \mathcal{K}$ satisfying $\int_{\Delta} \lambda_0 > 0$. Then every function $\hat{\lambda} \in \arg \min_{\lambda \in \mathcal{K}} f(\lambda)$ satisfies $\int_{\Delta} \hat{\lambda} = N$. Thus, the expected number of arrivals corresponding to the maximum likelihood estimator from \mathcal{K} equals the observed number of arrivals.*

Proof. Suppose $\hat{\lambda}$ is an optimal solution (implying $\int_{\Delta} \hat{\lambda} > 0$), and consider feasible solutions $c\hat{\lambda}$ with $c > 0$. We have $f(c\hat{\lambda}) = -c \int_{\Delta} \hat{\lambda} + N \ln c + \sum_{i \in \mathcal{I}} n_i \ln \int_{\Delta_i} \lambda$, and by assumption $\left. \frac{d}{dc} f(c\hat{\lambda}) \right|_{c=1} = 0$. The last equation gives $\int_{\Delta} \hat{\lambda} = N$, in which case $c = 1$ indeed maximizes $f(c\hat{\lambda})$. \square

Taking \mathcal{K} to be any cone of piecewise weighted-sum-of-squares polynomial splines the conditions of the Lemma are clearly satisfied.

Note that f is a concave function. Furthermore, if the partition $(\Delta_i)_i$ is a refinement of the

subdivision corresponding to the pieces of the sought spline estimator, then the objective function f is separable, leading to an optimization model of the form (4.4). (Except, of course, that we are maximizing a concave function instead of minimizing a convex one.)

Lemma 4.4 allows us to remove the first term of the objective function, provided we add the equation $\int_{\Delta} \lambda = N$ to our constraints. We shall not transform our model this way, as it would adversely affect the performance of the decomposition algorithm. (The constraint on the integral involves every decision variable.) Nevertheless, this lemma is necessary to meaningfully bound the feasible set of the decomposed problems (the sets \mathcal{X}_i with the notation of (4.4)). So far \mathcal{X}_i has been a weighted-sum-of-squares polynomial cone (nonnegative over a polyhedron), but by Lemma 4.4 we may restrict it (without loss of generality) to those weighted-sum-of-squares polynomials which integrate to at most N .

Numerical results. In our example $\Delta = [0, T] \times [0, X]$; $T = 1$ week, $X = 77.96$ miles. Considering the format of the data, the regions \mathcal{I} in the objective function can be rectangles no smaller than 1 minute by 0.1 miles, but even considerably larger rectangles are reasonable, given the rounding errors.

Figure 4.2 shows a biquadratic spline estimator with 28×13 pieces (so each piece corresponds to 6 hours and roughly 6 miles), the regions \mathcal{I} were 1 minute by 1 mile rectangles. The estimator was obtained using an AMPL [FGK02] implementation of Algorithm 1, in which the subproblems were solved by a (size-limited version of) KNITRO [Zie].

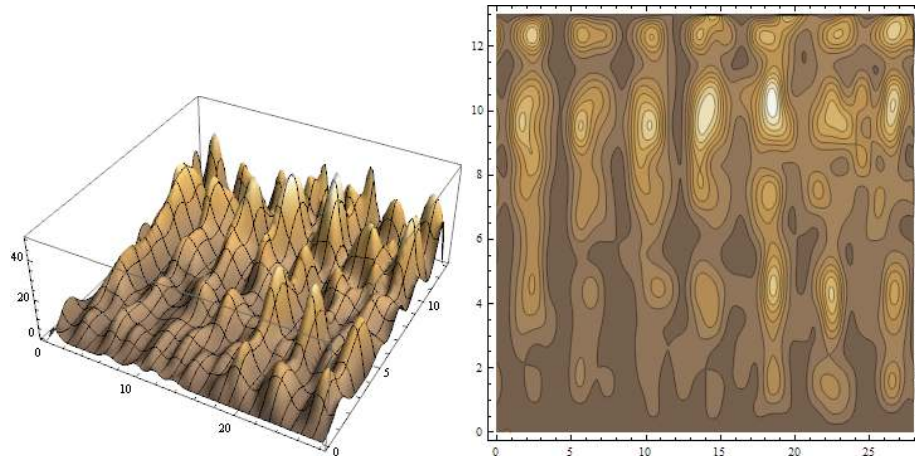


Figure 4.2: A piecewise biquadratic sum-of-squares spline estimator of the NJTP accident rate obtained using Algorithm 1. Left: three-dimensional plot. Right: contour plot.

4.3 Generalized sum-of-squares cones

In this section we generalize the results of [Nes00] on the semidefinite representability of sum-of-squares and weighted-sum-of-squares real-valued functions. We consider sums of squares of vectors in an arbitrary linear space equipped with a bilinear multiplication (mapping pairs of vectors to a possibly different linear space), and show how this serves as a semidefinite representable restriction of a number of intractable constraints arising in shape-constrained optimization, as well as in a number of other settings.

This section is primarily motivated by shape-constrained optimization problems, but the general theory we develop also has immediate applications in other areas. We give a brief outlook by showing a simple application to combinatorial optimization and polynomial programming (POP) in Section 4.3.3.

As the following examples show, not all shape constraints can be translated naturally to nonnegativity of some linear transform of the shape constrained real valued function.

1. (Convexity of a multivariate function.) A twice continuously differentiable real valued function f is convex over $S \subset \mathbb{R}^n$ if and only if its Hessian \mathbf{H} is positive semidefinite over S , which we denote by $\mathbf{H}(\mathbf{x}) \succcurlyeq 0, \forall \mathbf{x} \in S$. Magnani *et al.* [MLB05] consider this problem in the special case when f is a multivariate polynomial, and suggest the following approach: $\mathbf{H}(\mathbf{x}) \succcurlyeq 0$ for every $\mathbf{x} \in S$ if and only if $h(\mathbf{x}, \mathbf{y}) = \mathbf{y}^T \mathbf{H}(\mathbf{x}) \mathbf{y} \geq 0$ for every $\mathbf{x} \in S$ and $\mathbf{y} \in \mathbb{R}^n$. Since h is also a polynomial, the constraint is reduced to the nonnegativity of a multivariate polynomial. This reduction, however, is not entirely satisfying for two reasons. First, it calls for doubling the number of variables at the outset. Second, if f is not a polynomial, and \mathbf{H} is not a polynomial matrix, then \mathbf{H} is still a linear transform of f , however, $h(\mathbf{x}, \mathbf{y})$ may belong to an entirely different functional system, in which it is generally difficult to establish a connection between nonnegativity and being sum-of-squares.
2. (Covariance matrix estimation.) The above problems become even more evident in the case of covariance matrix estimation. Constraints of the form $\mathbf{P}(t) \succcurlyeq 0$ for every $t \in \mathbb{R}$ or for every $t \in [a, b]$ are semidefinite representable for polynomial matrices \mathbf{P} . (Over the entire real line this follows from a theorem of Youla, see Proposition 4.19, and for intervals we shall prove this later in this section, in Theorem 4.21.) On the other hand, the previous approach turns these constraints into nonnegativity of bivariate polynomials, which in general is not tractable.
3. (Functions with bounded curvature.) Consider a problem in which we are trying to find a twice

differentiable curve given by its parametric representation $\mathbf{x}(s) \in \mathbb{R}^n$, where $s \in [0, S]$ is the arc-length parameter, under the constraint that the curvature of the curve must be bounded above by some constant $C \geq 0$. This constraint can be written as

$$\|\mathbf{x}''(s)\| \leq C \quad \forall s \in (0, S),$$

where \mathbf{x}'' is the component-wise second derivative of the vector-valued function \mathbf{x} , and $\|\cdot\|$ is the Euclidean norm [Str88, Section 1-4]. Equivalently, the constraint can be written as

$$\begin{pmatrix} C \\ \mathbf{x}''(s) \end{pmatrix} \in \mathcal{Q}_{n+1} \quad \forall s \in (0, S),$$

where \mathcal{Q}_{n+1} is the $(n+1)$ -dimensional second-order cone, or the Lorentz cone [AG03].

The examples suggest that the notion of nonnegativity should be extended beyond real numbers to more general spaces. One such extension to higher dimensional linear spaces uses the partial order induced by a *proper cone*, that is a closed, pointed, convex and full-dimensional cone \mathcal{K} . (See the Appendix for details.) In this case our constraints are of the form $f(\mathbf{x}) \in \mathcal{K}, \forall \mathbf{x} \in S$, where f is a (perhaps multivariate) vector-valued function. Such requirements are called *cone-nonnegativity* or \mathcal{K} -*nonnegativity* constraints. This will generally be an intractable constraint, but we can try to find a tractable approximation for it in the form of f being “sum-of-squares” with respect to some multiplication of vector-valued functions. This is particularly appealing when \mathcal{K} is a symmetric cone, as symmetric cones are cones of squares with respect to a Euclidean (or, equivalently, formally real) Jordan algebra multiplication [FK94]. Kojima and Muramatsu [KM07] consider this problem in the special case when f is a vector-valued polynomial whose coefficients are multiplied according to a Euclidean Jordan algebra multiplication, and derive a semidefinite programming characterization for these sum-of-squares polynomials. Multivariate sum-of-squares matrix polynomials are also considered in [Koj03].

In this section we consider the following, even more general problem. We take an arbitrary bilinear mapping $\diamond: A \times A \rightarrow B$, where A and B are fixed finite-dimensional real linear spaces, and show that the set of vectors that are sums of squares of vectors from A (with respect to the multiplication \diamond) is a linear image of positive semidefinite matrices. The linear transformation is explicitly constructed using a polynomial number of multiplications. This generalizes the above-mentioned results from [Nes00] and [KM07].

Since in this section the objects we handle are all vectors from a number of different linear spaces

and algebras, we shall not typeset vectors and matrices in bold face, as customary in linear algebra. Instead, only in this section, all scalars, vectors, and matrices from abstract spaces and algebras are typeset italic (lightface), and boldface characters are reserved for tuples of vectors.

4.3.1 Semidefinite characterization of sums of squares

Consider two finite-dimensional real linear spaces A and B , and a bilinear mapping $\diamond: A \times A \rightarrow B$. We define the cone of *sum-of-squares* vectors $\Sigma \subseteq B$ by

$$\Sigma_{\diamond} \stackrel{\text{def}}{=} \left\{ \sum_{i=1}^N a_i \diamond a_i \mid N \geq 1; a_1, \dots, a_N \in A \right\}.$$

Alternatively, we may say that a Σ_{\diamond} defined this way is a *sum-of-squares cone*, or *SOS cone* for short. Clearly, Σ_{\diamond} is a convex cone, but it not necessarily proper, as it may have an empty interior, and it may contain lines. We will drop the subscript \diamond and simply write Σ if the multiplication in question is understood from the context. Note that the extreme rays of Σ are among perfect squares $a_i \diamond a_i$. Therefore, by the conic version of Carathéodory's theorem [Roc70, Corollary 17.1.2], each element of Σ can be written as sum of at most $\dim(B)$ squares.

Let $E = \{e_1, \dots, e_m\}$ be a basis of A , and $V = \{v_1, \dots, v_n\}$ be a basis of a vector space $B' \supseteq B$. Furthermore, let the vector $\lambda_{ij} \in \mathbb{R}^n$ ($i, j = 1, \dots, m$) be the coefficient vector of $e_i \diamond e_j \in B$ in the basis V :

$$e_i \diamond e_j = \sum_{\ell} \lambda_{ij\ell} v_{\ell}. \quad (4.8)$$

Finally, we define the linear operator $\Lambda: \mathbb{R}^n \rightarrow \mathbb{R}^{m \times m}$, coordinate-wise, by the formula

$$(\Lambda(w))_{ij} \stackrel{\text{def}}{=} \langle w, \lambda_{ij} \rangle \quad \forall w \in \mathbb{R}^n \quad (i, j = 1, \dots, m); \quad (4.9)$$

its adjoint operator is denoted by Λ^* . If \diamond is commutative, then Λ attains only symmetric values, and it is more natural to define Λ as an $\mathbb{R}^n \rightarrow \mathbb{S}^m$ operator, where \mathbb{S}^m is the space of $m \times m$ real symmetric matrices.

Our main theorem is the characterization of the sum-of-squares cone Σ as a linear image of the cone of positive semidefinite matrices. Such sets are called semidefinite representable (see also Definition A.9 in the Appendix).

Theorem 4.5. *Let $u = (u_1, \dots, u_n)^T \in \mathbb{R}^n$ be arbitrary. Then $\sum_{\ell=1}^n u_{\ell} v_{\ell} \in \Sigma$ if and only if there exists a real symmetric positive semidefinite matrix $Y \in \mathbb{R}^{m \times m}$ satisfying $u = \Lambda^*(Y)$.*

Proof. Let us assume first that $\sum_{\ell=1}^n u_\ell v_\ell \in \Sigma$, that is,

$$\sum_{\ell=1}^n u_\ell v_\ell = \sum_{k=1}^N a_k \diamond a_k$$

for some $a_k \in A$. Each a_k can be written in the basis E as $a_k = \sum_{j=1}^m y_j^{(k)} e_j$, with $y_j^{(k)} \in \mathbb{R}$. By choosing $Y = \sum_{k=1}^N y^{(k)} y^{(k)\top}$, which is clearly positive semidefinite, we obtain

$$\begin{aligned} \sum_{\ell=1}^n u_\ell v_\ell &= \sum_{k=1}^N a_k \diamond a_k = \sum_{k=1}^N \left(\left(\sum_{i=1}^m y_i^{(k)} e_i \right) \diamond \left(\sum_{j=1}^m y_j^{(k)} e_j \right) \right) \\ &= \sum_{k=1}^N \sum_{i=1}^m \sum_{j=1}^m y_i^{(k)} y_j^{(k)} (e_i \diamond e_j) \\ &= \sum_{i=1}^m \sum_{j=1}^m Y_{ij} (e_i \diamond e_j) = \sum_{i=1}^m \sum_{j=1}^m Y_{ij} \sum_{\ell=1}^n \lambda_{ij\ell} v_\ell, \end{aligned}$$

where the last equation comes from Equation (4.8). Using the symbol 1_ℓ for the ℓ -th unit vector $(0, \dots, 0, 1, 0, \dots, 0) \in \mathbb{R}^n$, the last expression can be simplified further:

$$\sum_{i=1}^m \sum_{j=1}^m Y_{ij} \sum_{\ell=1}^n \lambda_{ij\ell} v_\ell = \sum_{\ell=1}^n \langle Y, \Lambda(1_\ell) \rangle v_\ell = \sum_{\ell=1}^n \langle \Lambda^*(Y), 1_\ell \rangle v_\ell.$$

Since $\{v_\ell\}$ is a basis, comparing the coefficients in $\sum_{\ell} u_\ell v_\ell$ to those in the last expression yields $u = \Lambda^*(Y)$, completing the proof in the “if” direction.

To prove the converse claim, reverse the steps of the above proof: if $u = \Lambda^*(Y)$, and Y is positive semidefinite, then obtain vectors $y^{(k)}$ satisfying $Y = \sum_k y^{(k)} y^{(k)\top}$ (for example, from the spectral decomposition of Y), and use the above identities to deduce

$$\sum_{\ell=1}^n u_\ell v_\ell = \sum_k \left(\left(\sum_{i=1}^m y_i^{(k)} e_i \right) \diamond \left(\sum_{j=1}^m y_j^{(k)} e_j \right) \right) \in \Sigma. \quad \square$$

After fixing the basis $V = \{v_1, \dots, v_n\}$ for B' , the cone Σ is naturally identified with the cone $\{u \mid \sum_{\ell=1}^n u_\ell v_\ell \in \Sigma\}$. It shall raise no ambiguities that we denote the latter cone by Σ , too.

In optimization applications it is necessary to characterize the dual cone of Σ , denoted by Σ^* . Using our main theorem, and the well-known fact that the cone of positive semidefinite matrices is self-dual (in the space of symmetric matrices), Σ^* is easily characterized, especially when \diamond is commutative.

Theorem 4.6. *Using the notation above,*

$$\Sigma^* = \{v \mid \exists S \succcurlyeq 0, A = -A^T: \Lambda(v) = S + A\}.$$

In particular, if \diamond is commutative, then

$$\Sigma^* = \{v \mid \Lambda(v) \succcurlyeq 0\}.$$

Proof. By definition, a vector $v \in \mathbb{R}^n$ is in Σ^* if and only if $\langle u, v \rangle \geq 0$ for every $u \in \Sigma$. By Theorem 4.5, $u \in \Sigma$ if and only if $u = \Lambda^*(Y)$ for some $Y \succcurlyeq 0$, consequently $v \in \Sigma^*$ if and only if

$$\langle u, v \rangle = \langle \Lambda^*(Y), v \rangle = \langle Y, \Lambda(v) \rangle \geq 0 \quad (4.10)$$

for every $Y \succcurlyeq 0$. The dual cone of positive semidefinite matrices (embedded in $\mathbb{R}^{n \times n}$ as opposed to \mathbb{S}_n) is the cone of matrices that can be written in the form of $S + A$, where S is positive semidefinite and A is skew-symmetric. This proves our first claim.

If \diamond is commutative, then $\Lambda(v)$ is symmetric for every v . Hence, in the decomposition $\Lambda(v) = S + A$ we must have $A = 0$. This proves the second part of the theorem. \square

Without loss of generality we can assume that \diamond is commutative, because the sum-of-squares cone Σ_{\diamond_2} with respect to the multiplication \diamond_2 defined by $x \diamond_2 y = \frac{1}{2}(x \diamond y + y \diamond x)$ is identical to Σ_{\diamond} , and \diamond_2 is commutative.

The definition of the cone Σ is applicable to *every* bilinear operation. Not all operations result in interesting cones, however. We start with to trivial examples to see what the issues are.

Example 4.7. Let $A = B = B' = \mathbb{C}$, the algebra of complex numbers, viewed as a two-dimensional space over \mathbb{R} , with \diamond being the usual multiplication: $\begin{pmatrix} a_1 \\ b_1 \end{pmatrix} \diamond \begin{pmatrix} a_2 \\ b_2 \end{pmatrix} = \begin{pmatrix} a_1 a_2 - b_1 b_2 \\ a_1 b_2 + a_2 b_1 \end{pmatrix}$. Using the standard basis as V , we obtain that $\Lambda(w) = \begin{pmatrix} w_1 & w_2 \\ w_2 & -w_1 \end{pmatrix}$, and that

$$\Sigma = \left\{ \begin{pmatrix} Y_{11} & -Y_{22} \\ 2Y_{12} & \end{pmatrix} \mid \begin{pmatrix} Y_{11} & Y_{12} \\ Y_{12} & Y_{22} \end{pmatrix} \succcurlyeq 0 \right\} = \mathbb{R}^2$$

is the entire space. We conclude that every complex number is sum-of-squares, in concordance with the fact that every complex number is, in fact, a square. Similar situation holds for quaternions and octonions.

Example 4.8. Consider any finite-dimensional anticommutative algebra (\mathbb{A}, \diamond) over \mathbb{R} , that is,

assume $x \diamond y = -y \diamond x$ for every $x, y \in \mathbb{A}$. Each vector a_i in such algebras satisfies $a_i \diamond a_i = 0$, and it follows that $\Sigma = \{0\}$, and $\Sigma^* = \mathbb{A}$.

To avoid trivial examples such as the first one, it is useful to consider when Σ is a pointed cone. A convex cone K is *pointed* if it does not contain a line, or equivalently, if $0 \neq x \in K$ implies $-x \notin K$. As the following lemma shows, a condition sufficient to obtain a pointed Σ is that the multiplication \diamond be formally real: \diamond is said to be *formally real* if for every $a_1, \dots, a_k \in A$, $\sum_{i=1}^k (a_i \diamond a_i) = 0$ implies that each $a_i = 0$.

Lemma 4.9.

1. If \diamond is formally real, then Σ_\diamond is pointed.
2. Conversely, if Σ_\diamond is pointed and there are no nonzero nilpotent elements of degree two, then \diamond is formally real.

Proof. First suppose that \diamond is formally real and that for some vector x both x and $-x$ are in Σ_\diamond . Then $0 = x + (-x) \in \Sigma_\diamond$ is sum-of-squares: $0 = \underbrace{\sum_i a_i \diamond a_i}_x + \underbrace{\sum_i b_i \diamond b_i}_{-x}$, implying that each a_i and b_i is zero. Consequently $x = 0$ confirming that Σ_\diamond is pointed.

Conversely, if Σ_\diamond is pointed and $\sum_i a_i \diamond a_i = 0$ for some a_i , then each $a_i \diamond a_i = 0$, and thus $a_i = 0$ by assumption. This proves that \diamond is formally real. \square

Another special case of Theorem 4.5 is Nesterov's well-known characterization of sum-of-squares functional systems [Nes00].

Example 4.10. Let f_1, \dots, f_n be arbitrary functions mapping a set Δ to \mathbb{R} , and $A = \text{span}(\{f_1, \dots, f_n\})$. Then $\Sigma \subseteq \text{span}(\{f_1^2, f_1 f_2, \dots, f_n^2\})$. The semidefinite characterization of Σ obtained from Theorem 4.5 is identical to the one in [Nes00].

Example 4.11. As a special case of the previous example we can assume A to be the set of k -variate polynomials of degree at most n , and B the set of k -variate polynomials of degree at most $2n$. In that case $\Sigma = \Sigma_{2n,k}$ is the cone of SOS k -variate polynomials of degree $2n$.

Example 4.12. As a further specialization, let A be the set of univariate polynomials of degree n , and B the set of univariate polynomials of degree $2n$, with \diamond representing multiplication of polynomials. Representing the polynomials by their coefficient vectors in the standard basis $\{1, t, \dots, t^n\}$ for A and $\{1, t, \dots, t^{2n}\}$ for B , $\Lambda(w)$ is the Hankel matrix corresponding to w , that is $(\Lambda(w))_{ij} = w_{i+j}$. Note that this proves the last part of Proposition 2.3.

If A' is a proper subspace of A , then the sum-of-squares cone Σ'_\diamond corresponding to the restriction of \diamond to A' is a subset of Σ_\diamond , but it need not be a proper subset. If it is not, the cones Σ_\diamond and Σ_\diamond^* can be represented by smaller semidefinite matrices. A non-trivial example is given by the Cracovian algebra [Koc04].

Example 4.13. (Cracovian algebra.) Let $A = \mathbb{R}^{r \times k}$ and $B = \mathbb{R}^{r \times r}$, then the product $M \diamond N = MN^T$ is formally real. In this case it is obvious that Σ_\diamond is identical to the cone of positive semidefinite (real, symmetric) $r \times r$ matrices. If we follow the process of Theorem 4.6 for the Cracovian multiplication, then $\Lambda(w)$ will be represented by an $nr \times nr$ matrix; but we know that $r \times r$ matrices should suffice in this case. We obtain a representation with $r \times r$ matrices if we consider a restriction of \diamond to \mathbb{R}^r . In this case $\Sigma' = \Sigma$ the set of symmetric positive semidefinite matrices and the $\Lambda'(w)$ is an $r \times r$ matrix rather than $rk \times rk$.

Weighted-sum-of-squares cones

The above theory extends to the semidefinite representability of *weighted-sum-of-squares* vectors. Before we begin, let us recall the notion of a left multiplication operator.

Definition 4.14. Given a bilinear operation $\diamond : A \times A \rightarrow B$, for each element $x \in A$ there is a linear operator $L_\diamond(x)$ that maps every $y \in A$ to $x \diamond y$. We call $L_\diamond(x)$ the *left multiplication operator* of x .

We shall drop the subscript \diamond when the multiplication is clear from the context. This operator can be represented by an $n \times m$ matrix (also denoted by $L_\diamond(x)$ or L), where $m = \dim(A)$ and $n = \dim(B)$.

Lemma 4.15. Let Σ_\diamond be the SOS cone corresponding to $\diamond : A \times A \rightarrow B$, and let $w \in B$. In addition, let there be another bilinear operation $\circ : B \times B \rightarrow C$. Then the cone $\Sigma^w \stackrel{\text{def}}{=} w \circ \Sigma_\diamond$ is an SOS cone. Moreover, if $L_\circ(w)$ is injective, then Σ^w is linearly isomorphic to Σ_\diamond (that is, there exists a linear bijection between Σ^w and Σ_\diamond).

Proof. For the first claim, observe that the multiplication $\star : A \times A \rightarrow C$ defined by $x \star y \stackrel{\text{def}}{=} w \circ (x \diamond y)$ is linear in both x and y , and $\Sigma_\star = \Sigma^w$. If $L_\circ(w)$ is injective, it defines a bijective linear transformation from B to $L_\circ(w)(B)$, which is a subspace of C containing Σ^w . Thus, restriction of $L_\circ(w)$ to a linear transformation from B to $L_\circ(w)(B)$ is a bijection; applying this bijection to Σ_\diamond gives Σ^w proving their linear isomorphism. \square

Finally, we are ready define the (generalized) weighted-sum-of-squares cones.

Definition 4.16. Let $\mathbf{w} = (w_1, \dots, w_r) \in B^r$, and let $\Sigma^{w_i} = w_i \circ_i \Sigma_{\circ_i}$, where $\circ_i : A \times A \rightarrow B$ and $\circ_i : B \times B \rightarrow C$, for some finite-dimensional real linear spaces A , B , and C . Then the *weighted-sum-of-squares cone* $\Sigma^{\mathbf{w}}$ corresponding to A with *weights* w_i is

$$\Sigma^{\mathbf{w}} \stackrel{\text{def}}{=} \Sigma^{w_1} + \dots + \Sigma^{w_r}, \quad (4.11)$$

where $+$ denotes the Minkowski sum of sets.

Since the Minkowski sum of semidefinite representable sets is also semidefinite representable (this comes from the definition; for details see the Appendix), we find that the weighted-sum-of-squares cones are also semidefinite representable. The explicit semidefinite representation is also easy to construct, essentially identically to the proof of Theorem 4.5:

Theorem 4.17. *The weighted-sum-of-squares cone $\Sigma^{\mathbf{w}}$ defined in (4.11) is semidefinite representable. Specifically, there are r linear operators $\Lambda^k : \mathbb{R}^{(\dim A) \times (\dim A)} \rightarrow C$, $k = 1, \dots, r$, such that $u \in C$ if and only if there exists real symmetric positive semidefinite matrices $Y^{(1)}, \dots, Y^{(r)} \in \mathbb{R}^{m \times m}$ satisfying $u = \sum_{k=1}^r (\Lambda^k)^*(Y^{(k)})$.*

Proof. The proof is driven by the same idea as the proof of Theorem 4.5, and for computational applications it is particularly helpful to develop it in a basis-dependent manner, by explicitly constructing the linear and semidefinite constraints in the semidefinite representation.

In addition to the notation introduced for (4.11), let $\{e_1, \dots, e_m\}$ be a basis for A , $\{f_1, \dots, f_n\}$ be a basis of C , and let $\lambda_{ij}^k \in \mathbb{R}^m$ be the coefficient vector of $w_k \circ_k (e_i \diamond_k e_j)$:

$$w_k \circ_k (e_i \diamond_k e_j) = \sum_{\ell} \lambda_{ij\ell}^k f_{\ell},$$

and define Λ^k by the formula

$$(\Lambda^k(\cdot))_{ij} \stackrel{\text{def}}{=} \langle \cdot, \lambda_{ij}^k \rangle \quad (i, j = 1, \dots, m; k = 1, \dots, r).$$

Expressing all vectors from A in the above basis, and collecting terms analogously to the proof of Theorem 4.5, we obtain the following chain of identities:

$$\sum_{\ell=1}^n u_{\ell} f_{\ell} = \sum_{k=1}^r \left(w_k \diamond \left(\sum_{\ell=1}^{N_k} (a_{\ell}^k \diamond a_{\ell}^k) \right) \right) =$$

$$\begin{aligned}
&= \sum_{k=1}^r \left(w_k \diamond \left(\sum_{\ell=1}^{N_k} \left(\sum_{j=1}^m y_{\ell j}^{(k)} e_j \right) \diamond \left(\sum_{j=1}^m y_{\ell j}^{(k)} e_j \right) \right) \right) \\
&= \sum_{k=1}^r \left(w_k \diamond \left(\sum_{\ell=1}^{N_k} \sum_{i=1}^m \sum_{j=1}^m y_{\ell i}^{(k)} y_{\ell j}^{(k)} (e_i \diamond e_j) \right) \right) \\
&= \sum_{k=1}^r \sum_{\ell=1}^{N_k} \sum_{i=1}^m \sum_{j=1}^m y_{\ell i}^{(k)} y_{\ell j}^{(k)} (w_k \diamond (e_i \diamond e_j)) \\
&= \sum_{k=1}^r \sum_{i=1}^m \sum_{j=1}^m Y_{ij}^{(k)} (w_k \diamond (e_i \diamond e_j)) \\
&= \sum_{k=1}^r \sum_{i=1}^m \sum_{j=1}^m Y_{ij}^{(k)} \sum_{\ell=1}^n \lambda_{ij\ell}^{(k)} f_\ell \\
&= \sum_{k=1}^r \sum_{\ell=1}^n \langle Y^{(k)}, \Lambda^k(\mathbf{1}_\ell) \rangle f_\ell \\
&= \sum_{\ell=1}^n \sum_{k=1}^r \langle (\Lambda^k)^*(Y^{(k)}), \mathbf{1}_\ell \rangle f_\ell.
\end{aligned}$$

Comparing the coefficients of the first and last expression yields

$$u = \sum_{k=1}^r (\Lambda^k)^*(Y^{(k)}). \quad \square$$

Nesterov [Nes00] considers the weighted functional systems $\Sigma^{\mathbf{w}} = w_1(x)\Sigma_1 + \dots + w_k(x)\Sigma_k$, where each Σ_i is the SOS cone corresponding to the ordinary multiplication in the space spanned by a set of linearly independent functions $f_{1i}, \dots, f_{m_i i}$, all defined on a set Δ . First, for a single $w(x)\Sigma$ he sets $A_w = \text{span}\{\sqrt{w(x)}f_1(x), \dots, \sqrt{w(x)}f_m(x)\}$. Then clearly, the square functional system induced by A_w is $w(x)\Sigma$, where Σ is the SOS cone corresponding to $A = \text{span}\{f_1, \dots, f_m\}$. He then uses the SD-representability of the Minkowski sum of SD-representable sets to show that $w_1\Sigma_1 + \dots + w_l\Sigma_l$ is also SD-representable. This approach assumes that each $w_i(x)$ is nonnegative on Δ . Our development does not assume that we have access to “square roots” of w_i , nor does it rely on any notion of “nonnegativity”.

4.3.2 Application to vector-valued functions

While many of the examples in the previous sections can be regarded as basic, they serve as building blocks of some truly non-trivial results, such as the semidefinite characterization of (weighted) sums of squares of vector-valued functions and positive semidefinite matrix polynomials. All the results

of this section could be obtained by the direct application of Theorem 4.5; however, the following lemma considerably simplifies their presentation.

Recall that the tensor product of spaces $A_1 = \text{span}(\{e_i\})$ and $A_2 = \text{span}(\{f_j\})$ is defined as $\text{span}(\{e_i \otimes f_j\})$, where \otimes denotes the Kronecker product. As before, let A_i and B_i , $i = 1, 2$ be finite-dimensional real linear spaces, and $\diamond_i: A_i \times A_i \rightarrow B_i$ be bilinear mappings. We define the tensor product $\diamond: (A_1 \otimes A_2) \times (A_1 \otimes A_2) \rightarrow (B_1 \otimes B_2)$ of \diamond_1 and \diamond_2 by first defining it for tensor products: for $u_1, u_2 \in A_1$ and $v_1, v_2 \in A_2$,

$$(u_1 \otimes v_1) \diamond (u_2 \otimes v_2) \stackrel{\text{def}}{=} (u_1 \diamond_1 u_2) \otimes (v_1 \diamond_2 v_2).$$

The definition of \diamond is then extended to all elements of $A_1 \otimes A_2$ by (the unique) bilinear extension. The semidefinite representation of Σ_\diamond can be obtained from those of Σ_{\diamond_i} using the following observation.

Lemma 4.18. *Using the notation of Theorem 4.5 and (4.9), for every $w_1 \in B_1$ and $w_2 \in B_2$ we have $\Lambda_\diamond(w_1 \otimes w_2) = \Lambda_{\diamond_1}(w_1) \otimes \Lambda_{\diamond_2}(w_2)$. Thus, $\Lambda_\diamond = \Lambda_{\diamond_1} \otimes \Lambda_{\diamond_2}$.*

Proof. As above, let $\{e_i\}$ be a basis of A_1 and $\{f_i\}$ be a basis of A_2 . Then

$$\begin{aligned} (\Lambda_\diamond(w_1 \otimes w_2))_{(i_1, i_2), (j_1, j_2)} &= \langle w_1 \otimes w_2, (e_{i_1} \otimes f_{i_2}) \diamond (e_{j_1} \otimes f_{j_2}) \rangle \\ &= \langle w_1 \otimes w_2, (e_{i_1} \diamond_1 e_{j_1}) \otimes (f_{i_2} \diamond_2 f_{j_2}) \rangle = \langle w_1, e_{i_1} \diamond_1 e_{j_1} \rangle \langle w_2, f_{i_2} \diamond_2 f_{j_2} \rangle \\ &= (\Lambda_{\diamond_1}(w_1))_{i_1, j_1} (\Lambda_{\diamond_2}(w_2))_{i_2, j_2} = (\Lambda_{\diamond_1}(w_1) \otimes \Lambda_{\diamond_2}(w_2))_{(i_1, i_2), (j_1, j_2)}. \quad \square \end{aligned}$$

As an application, we can simplify the characterization of SOS cones of vector-valued functions whose multiplication decomposes to the multiplication of the coefficients and the multiplication of the underlying functional basis. In particular, we can handle SOS cones of the form

$$\{g_1(x) \diamond g_1(x) + \cdots + g_n(x) \diamond g_n(x) \mid g_i(x) = p_{i1}f_1(x) + \cdots + p_{im}f_m(x), \text{ with } p_{ij} \in A\},$$

where the f_i are real valued functions, by considering separately the SOS cone corresponding to multiplication in A and the SOS cone corresponding to the multiplication of the f_i .

In functional cones of Example 4.10, the cone of nonnegative functions is in general a proper superset of the cone of sum-of-squares functions. We have seen that in some exceptional cases, the two cones coincide. The most prominent example is the set of nonnegative univariate polynomials. There is an important extension of this result to symmetric-matrix-valued univariate polynomials; it relies on a consequence of Youla's spectral factorization theorem [You61]. For completeness we recall

the specific corollary we use.

Proposition 4.19 ([You61, Theorem 2 and Corollary 2]). *Let P be an $m \times m$ real symmetric univariate polynomial matrix, and let r be the largest integer such that P has at least one minor of order r that does not vanish identically. Then there exists an $m \times r$ polynomial matrix Q satisfying the identity $P(t) = Q(t)Q(t)^T$ identically.*

Let us define

$$\mathcal{P}_n^{\mathbb{S}^m} = \{(P_0, P_1, \dots, P_n) \in (\mathbb{S}^m)^{n+1} \mid P_0 + P_1 t + \dots + P_n t^n \succcurlyeq 0 \text{ for all } t \in \mathbb{R}\} \quad (4.12)$$

$$\mathcal{P}_{n,[a,b]}^{\mathbb{S}^m} = \{(P_0, P_1, \dots, P_n) \in (\mathbb{S}^m)^{n+1} \mid P_0 + P_1 t + \dots + P_n t^n \succcurlyeq 0 \text{ for all } t \in [a, b]\} \quad (4.13)$$

Youla's theorem states that $\mathcal{P}_n^{\mathbb{S}^m}$ is the SOS cone with respect to the Cracovian multiplication of polynomial matrices. Applying Lemma 4.18 along with the semidefinite representation of non-negative polynomials and Youla's theorem yields a semidefinite characterization analogous to the characterization of complex semidefinite matrix polynomials in [GHND03]:

Theorem 4.20. *The $m \times m$ matrix polynomial $P(t) = \sum_{i=0}^{2n} P_i t^i$ is positive semidefinite for every $t \in \mathbb{R}$ if and only if there exists a positive semidefinite block matrix $Y \in \mathbb{R}^{(n+1)m \times (n+1)m}$ consisting of blocks Y_{ij} , $i, j = 0, \dots, n$ of order m such that $P_\ell = \sum_{i+j=\ell} Y_{ij}$ for each $\ell = 0, \dots, 2n$.*

Proof. By Proposition 4.19, $P(t)$ is positive semidefinite for every $t \in \mathbb{R}$ if and only if it is sum-of-squares with respect to the multiplication $R(t) \diamond S(t) = R(t)S(t)^T$. We have characterized separately the sum-of-squares cone of the coefficient matrices (with respect to the Cracovian multiplication) in Example 4.13 and the sum-of-squares cone of ordinary polynomials (with respect to the ordinary polynomial multiplication) in Example 4.12. Now, the theorem follows from Lemma 4.18 and Theorem 4.5. \square

Finally, we can generalize Proposition 2.4 to symmetric matrix-valued polynomials.

Theorem 4.21. *For every $n \geq 1, m \geq 1$ the cones $\mathcal{P}_{n,[a,b]}^{\mathbb{S}^m}$ and $\mathcal{P}_{n,[a,\infty]}^{\mathbb{S}^m}$ are weighted sum-of-squares cones, specifically*

$$\begin{aligned} \mathcal{P}_{n,[a,\infty]}^{\mathbb{S}^m} &= \mathcal{P}_n^{\mathbb{S}^m} + (t-a)\mathcal{P}_n^{\mathbb{S}^m}, \\ \mathcal{P}_{n,[a,b]}^{\mathbb{S}^m} &= (t-a)\mathcal{P}_{n-1}^{\mathbb{S}^m} + (b-t)\mathcal{P}_{n-1}^{\mathbb{S}^m} && \text{when } n \text{ is odd,} \\ \mathcal{P}_{n,[a,b]}^{\mathbb{S}^m} &= (t-a)(b-t)\mathcal{P}_{n-2}^{\mathbb{S}^m} + \mathcal{P}_n^{\mathbb{S}^m} && \text{when } n \text{ is even.} \end{aligned}$$

Proof. For simplicity we assume $(a, b) = (0, 1)$, the argument for general intervals is essentially identical. (The general results can also be obtained from the special cases by a change of variables.) We will use the notation $A^{\circ 2}$ for squaring a matrix A with respect to the Cracovian multiplication.

First, suppose $P \in \mathcal{P}_{n,[0,\infty]}^{\mathbb{S}^m}$, and consider the polynomial Q defined by $Q(t) = Q(t^2)$. By assumption, Q is positive semidefinite everywhere, hence $Q = \sum_i Q_i^{\circ 2}$ by Youla's theorem (Proposition 4.19). Grouping the terms of the Q_i by the parity of their degrees we write $Q_i(t) = R_i(t^2) + tS_i(t^2)$, and since Q has no odd degree terms,

$$P(t^2) = Q(t) = \sum_i (R_i(t^2))^{\circ 2} + t^2 \sum_i (S_i(t^2))^{\circ 2}.$$

Taking $R = \sum_i R_i^{\circ 2}$ and $S = \sum_i S_i^{\circ 2}$ we have $P(t^2) = R(t^2) + t^2 S(t^2)$, implying $P(t) = R(t) + tS(t)$ with sum-of-squares matrices R and S , as claimed. The bounds on the degrees of follow from the fact that the degree of each R_i is at most n .

Next, suppose $P \in \mathcal{P}_{n,[0,1]}^{\mathbb{S}^m}$, and consider the polynomial $Q(t) = (1+t)^n P(t/(1+t))$: by assumption it is positive semidefinite for all $t \geq 0$, and so by the first claim of our theorem, $Q(t) = R(t) + tS(t)$ identically for some $R = \sum_i R_i^2$ of degree at most n and $S = \sum_i S_i^2$ of degree at most $n-1$.

Observe that $P(t) = (1-t)^n Q(t/(1-t))$. If $n = 2k+1$, then this yields

$$\begin{aligned} P(t) &= (1-t)^{2k+1} \sum_i R_i^{\circ 2}(t/(1-t)) + (1-t)^{2k+1} \frac{t}{1-t} \sum_i S_i^{\circ 2}(t/(1-t)) \\ &= (1-t) \sum_i ((1-t)^k R_i(t/(1-t)))^{\circ 2} + t \sum_i ((1-t)^k S_i(t/(1-t)))^{\circ 2}. \end{aligned}$$

If $n = 2k$, then

$$\begin{aligned} P(t) &= (1-t)^{2k} \sum_i R_i^{\circ 2}(t/(1-t)) + (1-t)^{2k} \frac{t}{1-t} \sum_i S_i^{\circ 2}(t/(1-t)) \\ &= \sum_i ((1-t)^k R_i(t/(1-t)))^{\circ 2} + t(1-t) \sum_i ((1-t)^{k-1} S_i(t/(1-t)))^{\circ 2}. \end{aligned}$$

The degree bounds come from the degree bound in the first claim of the theorem. \square

Similarly, Lemma 4.18 can be used to characterize the sums of squares of *arbitrary* vector-valued functions, not only of semidefinite valued matrix polynomials. More precisely, let (\mathbb{A}, \circ) be a (not necessarily commutative or associative) finite-dimensional algebra, f_1, \dots, f_n be given real valued functions, and let A be the space

$$\left\{ \sum_{i=0}^n v_i f_i \mid v_i \in \mathbb{A} \right\}.$$

Then the cone $\Sigma = A^2$ is semidefinite representable via Theorem 4.5, and a representation can be conveniently constructed using Lemma 4.18. This generalizes Lemma 2 of [KM07], where a similar characterization is obtained in the special case when $f_i = t^i$, and (\mathbb{A}, \circ) is a Euclidean (or formally real) Jordan algebra.

4.3.3 Further applications

Smallest Enclosing Ellipsoids of Curves

The smallest enclosing ball (SEB) and the smallest enclosing ellipsoid (SEE) (also often called minimum enclosing ball and ellipsoid) problems ask for the sphere or ellipsoid of minimum volume that contains a finite set of given points. These problems have been thoroughly studied, partly because of their important applications in machine learning, see for example [Bur98, CVBM02, KMY03]. Both of them admit simple second-order cone programming (SOCP) and semidefinite programming (SDP) formulations: denoting the input points by $Y = \{y_1, \dots, y_m\}$, the smallest sphere, with center x and radius r , containing Y is determined by the SOCP

$$\text{minimize } r \quad \text{subject to } \|x - y_i\| \leq r, i = 1, \dots, m,$$

while the ellipsoid of smallest volume is determined by the SDP

$$\text{maximize } (\det A)^{(1/n)} \quad \text{subject to } A \succ 0; \quad \|x - Ay_i\| \leq 1, i = 1, \dots, m.$$

Equivalently, the objective function can be replaced by $\ln \det A$, which results in a convex optimization problem with an SD-representable feasible set and a convex (but not SD-representable) objective function.

We consider the following generalization of the SEB and SEE problems: given a closed parametric curve $p(t), t \in [a, b]$ find the sphere or ellipsoid of minimum volume that contains all points of the curve. Replacing the finite set of constraints involving y_i by the constraints involving $p(t)$ we obtain optimization problems with a continuum of constraints over the cone $\mathcal{Q}_{n+1} = \{(x_0, \bar{x}) : x_0 \in \mathbb{R}, \bar{x} \in \mathbb{R}^n, \text{ and } x_0 \geq \|\bar{x}\|\}$. It is well known that \mathcal{Q}_{n+1} is a symmetric cone; in fact it is the cone of squares of the Jordan algebra $(\mathbb{R}^{n+1}, \circ)$ with binary operation $(x_0, \bar{x}) \circ (y_0, \bar{y}) = (\langle \bar{x}, \bar{y} \rangle + x_0 y_0, x_0 \bar{y} + y_0 \bar{x})$. Now, if p belongs to any class of functions whose cone of nonnegative functions is linearly isomorphic to the cone of nonnegative polynomials (for examples see also Section 5.1.4), then the set $\{p \mid p(t) \in \mathcal{Q} \forall t \in \Delta\}$ is an SOS cone for appropriate intervals Δ , and thus it is SD-representable.

Example 4.22. Figure 4.3 shows a parametric curve

$$(x(t), y(t)) = \sum_{k=1}^n p_{2k-1} \sin(kt) + p_{2k} \cos(kt), \quad t \in [0, 2\pi]$$

with $n = 3$ and coefficient vectors p_1, \dots, p_6 that are independent random vectors chosen uniformly from the interval $[-1, 1]^2$. Since p is a trigonometric polynomial with coefficients in \mathbb{R}^{n+1} , by the preceding discussion the condition $\|x - Ap(t)\| \leq 1$ is SD-representable. The minimal circle and minimal ellipse containing $\{p(t) \mid t \in [0, 2\pi]\}$ is shown in Figure 4.3. Obviously this example can be extended to arbitrary dimensions, as well as to enclosing a (finite) collection of curves, without any difficulty.




Figure 4.3: A trigonometric polynomial curve (blue, solid) together with its smallest enclosing circle (green, dashed) and its smallest enclosing ellipsoid (red, dot-dashed).

Polynomial Programming and Combinatorial Optimization

For a very brief section we shall wander away from our central topic to illustrate how general Theorems 4.5 and 4.17 are, and why we hope that it will find applications in entirely different areas of optimization, other than optimization with shape constraints.

Polynomial multiplication with respect to a (fixed) polynomial ideal is a bilinear operation. The application of Theorem 4.5 to the space of polynomials of a given degree yields simplified SDP relaxations and hierarchies to polynomial programs with equality constraints. Consider the POP

$$\text{maximize } f(x) \quad \text{subject to } h_i(x) = 0, \quad i = 1, \dots, m, \quad (4.14)$$

where f and each h_i are given n -variate polynomials. An equivalent formulation is

$$\text{minimize } c \text{ subject to } f(x) - c \geq 0 \forall x \in \{x: h_i(x) = 0, i = 1, \dots, m\}.$$

This problem is NP-hard, but the following restriction may be tractable:

$$\text{minimize } c \text{ subject to } f(x) - c \in \Sigma, \tag{4.15}$$

where Σ is the sum-of-squares cone with respect to polynomial multiplication modulo the ideal I generated by $\{h_i\}$. Depending on $\{h_i\}$, the bottleneck in the solution of (4.15) may be the computation of Λ , which involves finding a Gröbner basis of I , and a large number of multivariate polynomial divisions.

This method can be very effective when h_i form a Gröbner basis, and division modulo I is simple, as in the following example.

Example 4.23. Consider a graph $G = (V, E)$ with $|V| = n$ vertices and $|E| = m$ edges. Letting $f(x) = \sum_{i=1}^n x_i$ and $h_{ij}(x) = x_i x_j, ij \in E$, the solution to the problem (4.14) is the stability number of G . A sequence of semidefinite relaxations to the stable set problem is obtained if we replace (4.14) by (4.15), and constrain $f - c$ to be a sum of squares of polynomials of a given total degree. (Higher degrees give tighter relaxations.) Theorem 4.6 gives the dual SDPs of these relaxations, which are identical to Laurent’s simplified characterization [Lau03] of the Lasserre hierarchy [Las01a] applied to the stable set problem. This also serves as a new, simpler, proof of [Lau03, Lemma 20].

4.4 Conclusion

We considered several aspects of shape-constrained spline estimation in the multivariate setting. We introduced piecewise weighted-sum-of-squares splines and gave a condition for them to approximate continuous functions uniformly as the mesh size tends to zero. We proposed a decomposition method based on a partial linearization of the augmented Lagrangian for optimizing nonnegative splines with a very large number of pieces. Some illustrative numerical examples were included. Finally, we introduced the notion of weighted-sum-of-squares cones in arbitrary real linear spaces, and showed that they are semidefinite representable. This result has applications in both univariate and multivariate shape-constrained optimization.

Chapter 5

Direct optimization methods for nonnegative polynomials

While the cones of positive polynomials \mathcal{P}_{2n+1} and $\mathcal{P}_{n+1}^{[a,b]}$ and their dual cones are semidefinite representable, translating optimization problems involving these cones to the corresponding semidefinite constraints is not completely satisfactory. Recall from Proposition 2.3 that the semidefinite characterization of a vector in \mathcal{P}_{2n+1} requires a general $(n+1) \times (n+1)$ symmetric matrix to represent a $(2n+1)$ -dimensional vector; hence this approach effectively squares the number of variables. A second argument against this approach is that the semidefinite characterization of the moment cones \mathcal{M}_{2n+1} and $\mathcal{M}_{n+1}^{[a,b]}$ involve positive semidefinite Hankel matrices, which are notoriously ill-conditioned [Tyr94, Bec97]. In this chapter we investigate how the two most common general approaches to conic optimization problems, barrier methods and homotopy methods, can be adapted to positive polynomials and moment cones.

Until very recently, the only approach to this problem in the existing literature has been the application of the universal barrier function of Nesterov and Nemirovskii [NN94] to the cone of nonnegative polynomials, or more generally, nonnegative functions in Chebyshev systems. However, computing the universal barrier function and its derivatives for these cones appears to be difficult. Faybusovich gives a simplified formula to evaluate the universal barrier in [Fay02], but this formula involves the evaluation of the Pfaffian of a matrix whose entries given by integrals of general rational functions. (Furthermore, evaluation of the derivatives of the universal barrier is not discussed.) Evaluating the Pfaffian (determinant) is computationally very expensive, and it cannot be considered a really practical method; see [FMT08] for numerical experiments with this algorithm. No similar

formula is known for moment cones.

In addition to barrier methods, another large family of interior-point algorithms are homotopy methods. These methods solve a system of equations derived from the complementary slackness conditions of the problem at hand. The equations from the optimality conditions are perturbed, and the perturbed system is solved (approximately) by some general purpose method, such as Newton's method. In each iteration the perturbation is made smaller, and the approximate solution to the previous system is used as an initial solution to find an approximate solution to the less perturbed system. After a number of such iterations the approximate solution of the slightly perturbed system approximates the optimal solution to the original optimization problem.

In Section 5.1 we follow this approach, and examine the complementary slackness conditions for the cones $\mathcal{P}_{n+1}^{[a,b]}$ and $\mathcal{M}_{n+1}^{[a,b]}$. However, the main result of this section is negative: we show that the complementary slackness conditions for these cones cannot be written in as simple a way as for the first orthant, the second-order cone, or the semidefinite cone. More precisely, we show that any form of the complementary slackness conditions contains at most two linearly independent bilinear equations. As simple corollaries we also derive analogous results for cones of nonnegative trigonometric polynomials and their duals. These complement the main result in [Rud09, Chapter 3], that the number of linearly independent bilinear optimality conditions for \mathcal{P}_{2n+1} is at most four.

In Section 5.2 we return to the evaluation of barrier functions. We show that a self-concordant barrier function for the moment cone $\mathcal{M}_{n+1}^{[a,b]}$ (derived from its semidefinite representation) together with its gradient and Hessian can be computed in $\mathcal{O}(n^2 \log n)$ time using fast Fourier transform. This allows us to optimize over nonnegative polynomials (or moments) using semidefinite programming as before, but without increasing the number of variables, and with substantial reduction in the running time. For example, this subroutine reduces the time of the short-step path-following method [Ren01] from $\mathcal{O}(n^6)$ to $\mathcal{O}(n^3)$ per iteration, while maintaining the iteration complexity of the algorithm.

Fast Fourier transform (or equivalently, fast polynomial multiplication) is available only for polynomials represented in specific bases. One such basis is the standard monomial basis; however, that is usually not the optimal choice for numerical computations. Another basis for which fast Fourier transform is available is the Chebyshev basis (of the first kind). We show that the gradient and the Hessian of the same barrier function can be computed with the same running time if the Chebyshev polynomial basis is used in place of the monomial basis.

5.1 Optimality conditions for nonnegative polynomials

Let $\mathcal{K} \subseteq \mathbb{R}^n$ be a proper convex cone, and consider the following set:

$$\mathcal{C}(\mathcal{K}) \stackrel{\text{def}}{=} \{(\mathbf{x}, \mathbf{s}) \mid \mathbf{x} \in \mathcal{K}, \mathbf{s} \in \mathcal{K}^*, \langle \mathbf{x}, \mathbf{s} \rangle = 0\}.$$

The set $\mathcal{C}(\mathcal{K})$ is called the *complementarity set* of the cone \mathcal{K} . From the complementary slackness conditions of conic optimization problems (see Proposition A.5 in the Appendix) we have that if \mathbf{x} and (\mathbf{y}, \mathbf{s}) are feasible solutions to the primal-dual pair

$$\begin{array}{ll} \text{maximize} & \langle \mathbf{c}, \mathbf{x} \rangle \\ \text{subject to} & \mathbf{A}\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \succ_{\mathcal{K}} \mathbf{0} \end{array} \qquad \begin{array}{ll} \text{maximize} & \langle \mathbf{b}, \mathbf{y} \rangle \\ \text{subject to} & \mathbf{c} - \mathbf{A}^T \mathbf{y} = \mathbf{s} \\ & \mathbf{s} \succ_{\mathcal{K}^*} \mathbf{0}, \end{array}$$

and $\langle \mathbf{x}, \mathbf{s} \rangle = 0$, then \mathbf{x} and (\mathbf{y}, \mathbf{s}) are optimal solutions to the primal and dual problem, respectively. Conversely, under certain regularity conditions, if either problem has a feasible solution, then both problems have optimal solutions \mathbf{x} and (\mathbf{y}, \mathbf{s}) satisfying $(\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K})$. Consequently, having a good characterization for $\mathcal{C}(\mathcal{K})$, at least in principle, may allow us to solve conic optimization problems by solving the feasibility problem

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \quad \mathbf{c} - \mathbf{A}^T \mathbf{y} = \mathbf{s}, \quad (\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K}).$$

Let us recall a few familiar examples.

Linear programming. In linear programming we have $\mathcal{K} = \mathcal{K}^* = \mathbb{R}_+^n$ (the first orthant), and $(\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K})$ if and only if $\mathbf{x} \geq \mathbf{0}$, $\mathbf{s} \geq \mathbf{0}$, and $x_i s_i = 0$ for $i = 1, \dots, n$. Note that the system

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \quad \mathbf{c} - \mathbf{A}^T \mathbf{y} = \mathbf{s}, \quad x_i s_i = 0, \quad i = 1, \dots, n$$

consists of as many, $2n + m$, independent equations as the number of variables. Finding a solution to this system by homotopy methods is the basic idea of a number of interior-point methods for linear programming.

Second-order cone programming. In second-order cone programming we have $\mathcal{K} = \mathcal{K}^* = \mathcal{Q}_{n+1}$ (the second-order cone or Lorentz cone), and using the notation $\mathbf{x} = (x_0, \dots, x_n)^T$ and $\mathbf{s} = (s_0, \dots, s_n)^T$,

the complementarity relation $(\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K})$ holds if and only if $\mathbf{x} \in \mathcal{Q}_{n+1}$, $\mathbf{s} \in \mathcal{Q}_{n+1}$, and

$$\begin{pmatrix} \langle \mathbf{x}, \mathbf{s} \rangle \\ x_0 \mathbf{s} + s_0 \mathbf{x} \end{pmatrix} = \begin{pmatrix} 0 \\ \mathbf{0} \end{pmatrix}.$$

As in the previous example, solving the system

$$\mathbf{Ax} = \mathbf{b}, \quad \mathbf{c} - \mathbf{A}^T \mathbf{y} = \mathbf{s}, \quad \begin{pmatrix} \langle \mathbf{x}, \mathbf{s} \rangle \\ x_0 \mathbf{s} + s_0 \mathbf{x} \end{pmatrix} = \begin{pmatrix} 0 \\ \mathbf{0} \end{pmatrix}$$

by homotopy methods is the basic idea in a number of interior-point methods for second-order cone programming.

Semidefinite programming. In semidefinite programming we have $\mathcal{K} = \mathcal{K}^* = \mathbb{S}_+^n$ (the cone of $n \times n$ positive semidefinite real symmetric matrices), and the complementarity relation $\langle \mathbf{X}, \mathbf{S} \rangle = \sum_{i,j} X_{ij} S_{ij} = 0$ holds if and only if $\mathbf{XS} + \mathbf{SX} = \mathbf{0}$. Hence, the heart of the optimality conditions for semidefinite programming is the system

$$\mathbf{Ax} = \mathbf{b}, \quad \mathbf{c} - \mathbf{A}^T \mathbf{y} = \mathbf{s}, \quad \mathbf{XS} + \mathbf{SX} = \mathbf{0}.$$

That such an approach might be fruitful for *every* proper convex cone \mathcal{K} is indicated by the following theorem.

Proposition 5.1 ([RNPA]). *For every proper cone $\mathcal{K} \subseteq \mathbb{R}^n$, $\mathcal{C}(\mathcal{K})$ is an n -dimensional manifold homeomorphic to \mathbb{R}^n .*

While this theorem is known in the conic optimization community, it appears that there is no written account on any proof of it. A proof, attributed to Osman Güler, is contained in the Appendix of the forthcoming paper [RNPA].

Intuitively this result is somewhat surprising even after the concrete examples above, since in a “generic” setting one would expect the complementarity set, the intersection of a $(2n - 1)$ -dimensional space and a $2n$ -dimensional cone, to be a $(2n - 1)$ -dimensional manifold. Instead, Proposition 5.1 suggests that there may be $n - 1$ independent valid equations for $\mathcal{C}(\mathcal{K})$ in addition to the complementarity condition for every proper convex cone \mathcal{K} , not only for symmetric cones.

In all of the above examples we can endow the underlying linear space V ($V = \mathbb{R}_+^n$, \mathcal{Q}_{n+1} , and \mathbb{S}_+^n , respectively) with a bilinear multiplication \circ , such that the in the definition of $\mathcal{C}(\mathcal{K})$ the

complementarity relation $\langle \mathbf{x}, \mathbf{s} \rangle = 0$ can be replaced by $\mathbf{x} \circ \mathbf{s} = \mathbf{0}$. In the above examples the algebras (V, \circ) are the Euclidean Jordan algebras [FK94, Koe99] which have been extensively studied and used in the design and analysis of (primal-dual) interior-point methods for optimization problems involving these cones [Fay97, AS00]. This (and the fact that establishing algebraic independence of general polynomial equations is extremely hard, but is easy for bilinear forms, which also appear in the optimality conditions of symmetric cones) motivated the further study of optimality conditions that are expressible *by bilinear forms alone* in the report [ANR05] and the thesis [Rud09]. To make the discussion on such optimality conditions precise, we need the following definitions.

Definition 5.2. Let $\mathcal{K} \in \mathbb{R}^n$ be a proper cone. The matrix $\mathbf{Q} \in \mathbb{R}^{n \times n}$ is a *bilinear complementarity relation* for \mathcal{K} if every $(\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K})$ satisfies $\mathbf{x}^T \mathbf{Q} \mathbf{s} = 0$.

Note that the bilinear complementarity relations for a cone form a linear space. Motivated by Proposition 5.1 and the examples in the beginning of this section, we are interested in those cones that possess at least n linearly independent bilinear complementarity conditions.

Definition 5.3. The *bilinearity rank* of a proper cone \mathcal{K} , denoted by $\beta(\mathcal{K})$, is the dimension of the space of its bilinear complementarity relations. A proper cone $\mathcal{K} \subseteq \mathbb{R}^n$ is said to be *bilinear* if $\beta(\mathcal{K}) \geq n$.

The main result in [ANR05] and [Rud09, Chapter 3] is the following.

Proposition 5.4 ([ANR05]). *The cones \mathcal{P}_{2n+1} and \mathcal{M}_{2n+1} are not bilinear, unless $n = 1$. Moreover, for every $n \geq 2$, $\beta(\mathcal{P}_{2n+1}) = \beta(\mathcal{M}_{2n+1}) = 4$.*

In this section we use techniques similar to those in [Rud09] to show that $\beta(\mathcal{P}_{n+1}^{[a,b]}) = \beta(\mathcal{M}_{n+1}^{[a,b]}) = 2$ for every positive integer n and reals $a < b$. We also show an analogous result for two types of trigonometric polynomials, and for “almost all” Müntz polynomials.

5.1.1 Preliminary observations

We start with three trivial observations.

Proposition 5.5. *For every proper cone \mathcal{K} , $\beta(\mathcal{K}) = \beta(\mathcal{K}^*)$.*

Proof. The complementarity sets of \mathcal{K} and \mathcal{K}^* are congruent: $(\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K})$ if and only if $(\mathbf{s}, \mathbf{x}) \in \mathcal{C}(\mathcal{K}^*)$. □

Proposition 5.6. *If two proper cones are linearly isomorphic (that is, if there is a non-singular linear transformation mapping one onto the other), then their bilinearity ranks are the same.*

Proof. Let \mathbf{A} be a nonsingular linear transformation such that $\mathbf{A}(\mathcal{K}_1) = \mathcal{K}_2$. Then $\mathcal{K}_2^* = \mathbf{A}^{-T}(\mathcal{K}_1^*)$. Furthermore, \mathbf{Q}_i ($i = 1, \dots, m$) are linearly independent bilinear complementarity conditions for \mathcal{K}_1 if and only if $\mathbf{A}^{-T}\mathbf{Q}_i\mathbf{A}^T$ ($i = 1, \dots, m$) define linearly independent bilinear complementarity conditions for \mathcal{K}_2 . \square

As an example, one can easily show that $\beta(\mathcal{P}_{n+1}^{[a,b]}) = \beta(\mathcal{P}_{n+1}^{[0,1]})$.

Corollary 5.7. *Every cone $\mathcal{P}_{n+1}^{[a,b]}$ is linearly isomorphic to $\mathcal{P}_{n+1}^{[0,1]}$. Hence, $\beta(\mathcal{P}_{n+1}^{[a,b]}) = \beta(\mathcal{P}_{n+1}^{[0,1]})$.*

Proof. The polynomial p is nonnegative over $[a, b]$ if and only if the polynomial q defined by $q(t) = p(\frac{t-a}{b-a})$ is nonnegative over $[0, 1]$. Furthermore, the coefficients of q can be obtained by a nonsingular linear transformation from the coefficients of p . \square

Our next observation suggests a method to bound the bilinearity rank of a cone from above in a constructive manner.

Proposition 5.8. *For every proper cone \mathcal{K} ,*

$$\beta(\mathcal{K}) = \text{codim}(\text{span}\{\mathbf{s}\mathbf{x}^T \mid (\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K})\}).$$

Proof. The matrix \mathbf{Q} is a bilinear complementarity condition for \mathcal{K} if and only if

$$0 = \mathbf{x}^T\mathbf{Q}\mathbf{s} = \langle \mathbf{s}\mathbf{x}^T, \mathbf{Q}^T \rangle \quad \forall (\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K}). \quad \square$$

Corollary 5.9. *Suppose $\mathcal{K} \subseteq \mathbb{R}^n$ is a proper cone. If there are k pairs of vectors $(\mathbf{x}_i, \mathbf{s}_i) \in \mathcal{C}(\mathcal{K})$, $i = 1, \dots, k$, such that the matrices $\mathbf{s}_i\mathbf{x}_i^T$ are linearly independent, then $\beta(\mathcal{K}) \leq n^2 - k$. In particular, if $k > n^2 - n$, then \mathcal{K} is not bilinear.*

This is the approach used in [RNPA] to determine the bilinearity rank of polyhedral cones:

Theorem 5.10 ([RNPA]). *Let \mathcal{K} be a polyhedral cone. Then $\beta(\mathcal{K}) = n$ if and only if \mathcal{K} is linearly isomorphic to \mathbb{R}_+^n . Otherwise $\beta(\mathcal{K}) = 1$.*

As it is not closely related to our discussion, we omit the proof.

In the upcoming proofs, we are not going to follow the method directly suggested by Corollary 5.9; instead, it will be easier to construct an *infinite family* of pairs $(\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K})$ which are then shown to contain k pairs that correspond to linearly independent matrices $\mathbf{s}_i\mathbf{x}_i^T$, $i = 1, \dots, k$. Considering that $(\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K})$ can only hold if $\mathbf{x} \in \text{bd}(\mathcal{K})$ and $\mathbf{s} \in \text{bd}(\mathcal{K}^*)$, we have the following corollary of Proposition 5.8.

Corollary 5.11 ([RNPA]). *Suppose $\mathcal{K} \subseteq \mathbb{R}^n$ is a proper cone. If*

$$\dim(\text{span}(\{\mathbf{x}\mathbf{s}^T \mid \mathbf{x} \in \text{bd}(\mathcal{K}), \mathbf{s} \in \text{bd}(\mathcal{K}^*)\})) \geq k,$$

then $\beta(\mathcal{K}) \leq n^2 - k$. In particular, if $k > n^2 - n$, then \mathcal{K} is not bilinear.

5.1.2 Nonnegative polynomials over a finite interval

Now we turn our attention to $\mathcal{P}_{n+1}^{[a,b]}$, and prove our main theorem of this section, which is $\beta(\mathcal{P}_{n+1}^{[a,b]}) = 2$ for every positive integer n and finite interval $[a, b]$. The lower bound is easy to establish, using the following observation.

Lemma 5.12. *Suppose \mathbf{Q} is a bilinear complementarity relation. Then the equation $\mathbf{x}^T \mathbf{Q} \mathbf{s} = 0$ is satisfied by every $(\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K})$ if and only if it is satisfied by every $(\mathbf{x}, \mathbf{s}) \in \mathcal{C}(\mathcal{K})$ such that \mathbf{x} is an extreme vector of \mathcal{K} and \mathbf{s} is an extreme vector of \mathcal{K}^* . The same is true if we replace “extreme” with “boundary”.*

Proof. The *only if* direction is obvious. To show the converse implication, recall that every $\mathbf{x} \in \mathcal{K}$ and $\mathbf{s} \in \mathcal{K}^*$ can be expressed as a sum of finitely many extreme (or boundary) vectors of \mathcal{K} and \mathcal{K}^* , respectively. Furthermore, if $\mathbf{x} = \sum_{i=1}^k \mathbf{x}_i$ and $\mathbf{s} = \sum_{j=1}^{\ell} \mathbf{s}_j$, then $\langle \mathbf{x}, \mathbf{s} \rangle = 0$ if and only if $\langle \mathbf{x}_i, \mathbf{s}_j \rangle = 0$ for every $1 \leq i \leq k$, $1 \leq j \leq \ell$. Therefore, if $\langle \mathbf{x}, \mathbf{s} \rangle = 0$, and the complementarity relation is satisfied by every orthogonal pair of extreme (boundary) vectors, then $\langle \mathbf{x}_i, \mathbf{s}_j \rangle = 0$ for every $1 \leq i \leq k$, $1 \leq j \leq \ell$, and

$$\mathbf{x}^T \mathbf{Q} \mathbf{s} = \left(\sum_{i=1}^k \mathbf{x}_i \right)^T \mathbf{Q} \left(\sum_{j=1}^{\ell} \mathbf{s}_j \right) = \sum_{i=1}^k \sum_{j=1}^{\ell} \mathbf{x}_i^T \mathbf{Q} \mathbf{s}_j = 0. \quad \square$$

Lemma 5.13. *For every integer $n \geq 1$ and real numbers $a < b$, $\beta(\mathcal{P}_{n+1}^{[a,b]}) \geq 2$.*

Proof. We show that the following bilinear complementarity relations are satisfied by every $(\mathbf{p}, \mathbf{c}) \in \mathcal{C}(\mathcal{P}_{n+1}^{[a,b]})$:

$$\sum_{i=0}^n p_i c_i = 0, \quad (5.1a)$$

$$\sum_{i=0}^{n-1} (n-i)p_i c_{i+1} - (a+b) \sum_{i=1}^{n-1} i p_i c_i - ab \sum_{i=1}^n i p_i c_{i-1} = 0. \quad (5.1b)$$

It is easy to see that these conditions are indeed linearly independent. The equation (5.1a) is simply the orthogonality condition, which always holds by the definition of $\mathcal{C}(\mathcal{P}_{n+1}^{[a,b]})$.

By Lemma 5.12 it is enough to show that (5.1b) is satisfied for pairs of vectors $(\mathbf{p}, \mathbf{c}) \in C(\mathcal{P}_{n+1}^{[a,b]})$ where \mathbf{c} is a boundary vector of $\mathcal{M}_{n+1}^{[a,b]}$.

Let $p \in \mathcal{P}_{n+1}^{[a,b]}$ be a polynomial, and \mathbf{p} its coefficient vector. By the definition of $\mathcal{M}_{n+1}^{[a,b]}$, every boundary vector of $\mathcal{M}_{n+1}^{[a,b]}$ is of the form $\mathbf{c} = \alpha \mathbf{c}_{n+1}(t_0)$ for some $\alpha > 0$ and $t_0 \in [a, b]$. (Recall the notation from Definition 2.16.) Furthermore, $\mathbf{c}_{n+1}(t_0)$ is orthogonal to \mathbf{p} if and only if $\sum_i p_i t^i = p(t_0) = 0$. Therefore, $(\mathbf{p}, \mathbf{c}) \in C(\mathcal{P}_{n+1}^{[a,b]})$ implies $p(t_0) = 0$.

Since $p(t) \geq 0$ for every $t \in [a, b]$, every root of p has even multiplicity, except possibly for a and b , and hence

$$(t_0 - a)(b - t_0)p'(t_0) = 0.$$

Finally, this last equality and $p(t_0) = 0$ together imply

$$nt_0p(t_0) + (t_0 - a)(b - t_0)p'(t_0) = 0,$$

equivalent to (5.1b), since

$$\begin{aligned} nt_0p(t_0) + (t_0 - a)(b - t_0)p'(t_0) &= \\ &= \sum_{i=0}^n np_i t_0^{i+1} + \sum_{i=1}^n (t_0 - a)(b - t_0)ip_i t_0^{i-1} \\ &= \sum_{i=0}^{n-1} (n - i)p_i t_0^{i+1} - (a + b) \sum_{i=0}^{n-1} ip_i t_0^i - ab \sum_{i=1}^n ip_i t_0^{i-1} \\ &= \sum_{i=0}^{n-1} (n - i)p_i c_{i+1} - (a + b) \sum_{i=1}^{n-1} ip_i c_i - ab \sum_{i=1}^n ip_i c_{i-1}. \quad \square \end{aligned}$$

Finding matching the upper bounds is more tedious. They can be obtained separately for polynomials of even and odd degree, and the different representations of the extreme rays suggest that a unified proof might be difficult. The key ideas of both proofs are the same, but the details of the odd degree case are considerably more involved.

We will need the following elementary fact from linear algebra.

Lemma 5.14. *Let k be a positive integer and let $\mathcal{B} = \{b_1, \dots, b_k\}$ be a set of linearly independent vectors in a real vector space. For a set $\{m_1, \dots, m_k\} \subset \text{span}(\mathcal{B})$ consider the coordinates $\alpha_{i,j} \in \mathbb{R}$ ($i, j = 1, \dots, k$) uniquely defined by the representations $m_i = \sum_{j=1}^k \alpha_{i,j} b_j$. (We refer to this as the*

\mathcal{B} -representation of m_i .) If the conditions

$$\begin{aligned}\alpha_{i,i} &\neq 0 && \text{for all } 1 \leq i \leq k, \\ \alpha_{i,j} &= 0 && \text{for all } 1 \leq i < j \leq k\end{aligned}$$

hold, then the set $\{m_1, \dots, m_k\}$ is also linearly independent.

Proof. Observe that the matrix $(\alpha_{i,j})_{k \times k}$ is lower triangular with a nonzero diagonal, so it is non-singular. \square

We are going to use the following, formally more general version of the above lemma:

Corollary 5.15. *Let $\mathcal{B} \subset \mathbb{R}[x_1, \dots, x_n]$ be a finite set of linearly independent polynomials and consider a set $\mathcal{M} \subset \text{span}(\mathcal{B})$ with coordinates $\alpha_{m,b}$ ($m \in \mathcal{M}$, $b \in \mathcal{B}$) defined by the representations $m = \sum_{b \in \mathcal{B}} \alpha_{m,b} b$. Assume that there exists an injection $\varphi : \mathcal{B} \rightarrow \mathcal{M}$ and a linear order \prec on $\varphi(\mathcal{B})$ such that*

$$\begin{aligned}\alpha_{\varphi(b),b} &\neq 0 && \text{for all } b \in \mathcal{B}, \\ \alpha_{\varphi(b),d} &= 0 && \text{for all } b, d \in \mathcal{B} \text{ satisfying } \varphi(b) \prec \varphi(d).\end{aligned}$$

Then $\dim(\text{span}(\mathcal{M}(\mathbb{R}^n))) = |\mathcal{B}|$, where $\mathcal{M}(\mathbb{R}^n) \stackrel{\text{def}}{=} \{(m(\mathbf{x}))_{m \in \mathcal{M}} \mid \mathbf{x} \in \mathbb{R}^n\}$.

Proof. Let $k = |\mathcal{B}|$. It is well-known that for a vector $P = (p_1, \dots, p_k) \in (\mathbb{R}[x_1, \dots, x_n])^k$ consisting of linearly independent polynomials we have $\dim(\text{span}(P(\mathbb{R}^n))) = k$, therefore it suffices to find a k -element linearly independent subset of \mathcal{M} . As φ is injective, there exists an indexing $\mathcal{B} = \{b_1, \dots, b_k\}$ such that $\varphi(b_1) \prec \dots \prec \varphi(b_k)$. Let $m_i = \varphi(b_i) \in \mathcal{M}$ (for all $i = 1, \dots, k$). It is easy to verify that the sets $\{b_1, \dots, b_k\}$ and $\{m_1, \dots, m_k\}$ satisfy the conditions of Lemma 5.14. Consequently the set $\{m_1, \dots, m_k\} \subset \mathcal{M}$ is linearly independent, which implies our claim. \square

Theorem 5.16. *The cone $\mathcal{P}_{2n+1}^{[a,b]}$ is not bilinear. More specifically, for every n , $\beta(\mathcal{P}_{2n+1}^{[a,b]}) = 2$.*

Proof. We have seen that $\beta(\mathcal{P}_{2n+1}^{[a,b]}) \geq 2$, only the \leq direction needs proof.

Consider the matrix valued functions $M : \mathbb{R}^{n+2} \mapsto \mathbb{R}^{(2n+1) \times (2n+1)}$ defined as

$$M(t_1, \dots, t_n; \alpha, \beta) = \mathbf{c}_{2n+1}(t_1) \mathbf{p}^T + \alpha \mathbf{c}_{2n+1}(a) \mathbf{p}_a^T + \beta \mathbf{c}_{2n+1}(b) \mathbf{p}_b^T,$$

where $\mathbf{p}, \mathbf{p}_a, \mathbf{p}_b \in \mathcal{P}_{2n+1}^{[a,b]}$ are the coefficient vectors of the polynomials $p(x) = \prod_{k=1}^n (x - t_k)^2$, $p_a(x) = x - a$, and $p_b(x) = b - x$, respectively. It is easy to verify that the entries of $M = (m_{i,j})_{i,j=0}^{2n}$

satisfy the polynomial equation

$$\sum_{j=0}^{2n} m_{i,j} x^j \equiv t_1^i \prod_{k=1}^n (x - t_k)^2 + \alpha a^i (x - a) + \beta b^i (b - x). \quad (5.2)$$

The polynomials $p(x)$, $p_a(x)$ and $p_b(x)$ are clearly nonnegative over $[a, b]$, and by Proposition 2.19, $\langle \mathbf{p}_a, \mathbf{c}_{2n+1}(a) \rangle = \langle \mathbf{p}_b, \mathbf{c}_{2n+1}(b) \rangle = \langle \mathbf{p}, \mathbf{c}_{2n+1}(t_1) \rangle = 0$. We need to show that $\dim(\text{span}(M(\mathbb{R}^{n+2}))) = (2n+1)^2 - 2$, then the theorem follows from Corollary 5.11. Finally, we will show this equality using the sufficient condition presented in Corollary 5.15, with the set $\{m_{i,j}\}$ playing the role of set \mathcal{M} .

For the rest of the proof, let us assume that $a = 0$, $b = 1$; Corollary 5.7 guarantees that this is without loss of generality. By convention, $a^0 = 1$ in the rest of the proof.

Let us define the n -variate polynomials $\Pi(k, \ell)$ by

$$\Pi(k, \ell)(t_1, \dots, t_n) \stackrel{\text{def}}{=} \sum_{\substack{0 \leq \alpha_2, \dots, \alpha_n \leq 2 \\ \alpha_2 + \dots + \alpha_n = \ell}} t_1^k \prod_{j=2}^n (-2)^{(\alpha_j \bmod 2)} t_j^{\alpha_j}, \quad (5.3)$$

whenever $0 \leq k \leq 2n+2$ and $0 \leq \ell \leq 2n-2$; for values of k and ℓ outside these ranges let us define $\Pi(k, \ell)$ to be the zero polynomial. Let \mathcal{M} be the set of entries of the matrix M , and let \mathcal{B} be the set

$$\mathcal{B} = \{\alpha, \beta\} \cup \{\Pi(k, \ell) \mid 0 \leq k \leq 2n+2, 0 \leq \ell \leq 2n-2\}.$$

The elements of \mathcal{B} are considered as polynomials of $n+2$ variables, $t_1, \dots, t_n, \alpha, \beta$. The set \mathcal{B} is linearly independent, because no two polynomials in it share a common monomial. It follows from these definitions that for every $0 \leq i, j \leq 2n$,

$$m_{i,j} = \Pi(i, 2n-j) + \Pi(i+1, 2n-1-j) + \Pi(i+2, 2n-2-j) + m'_{i,j}, \quad (5.4)$$

where

$$m'_{i,j} = \begin{cases} \beta & j = 0 \\ \alpha - \beta & i = 0, j = 1 \\ -\beta & i \geq 1, j = 1 \\ 0 & \text{otherwise} \end{cases}. \quad (5.5)$$

We now introduce an injection $\varphi: \mathcal{B} \mapsto \mathcal{M}$ by defining its inverse (where it exists): let $m_{i,j}$ be the

image of the polynomial

$$\varphi^{-1}(m_{i,j}) = q_{i,j} \stackrel{\text{def}}{=} \begin{cases} \Pi(i, 2n - j) & j \geq \max\{2, i\} \\ \Pi(i + 2, 2n - 2 - j) & j \leq \min\{i - 1, 2n - 2\} \\ \beta & i = 0, j = 0 \\ \alpha & i = 0, j = 1 \\ \text{not defined} & \text{otherwise} \end{cases}. \quad (5.6)$$

Note that, φ^{-1} assigns a polynomial to every entry $m_{i,j}$ of M except for $m_{1,1}$ and $m_{2n,2n-1}$, and it assigns different polynomials to different entries of M . Consequently, φ is indeed an injection, and Equations (5.4) and (5.5) show that the coefficient of $q_{i,j}$ in the \mathcal{B} -representation of $m_{i,j}$ is 1.

Let us define a linear order \succ on $\varphi(\mathcal{B})$ in the following way: $m_{i_1,j_1} \succ m_{i_2,j_2}$ precisely when one of the following four conditions holds:

1. $(i_1, j_1) = (0, 1)$;
2. $i_1 - j_1 \geq 1 > i_2 - j_2$;
3. $i_1 - j_1 \geq 1, i_2 - j_2 \geq 1$, and either $i_1 > i_2$, or $i_1 = i_2$ but $j_1 < j_2$;
4. $i_1 - j_1 < 1, i_2 - j_2 < 1$, and either $j_1 < j_2$, or $j_1 = j_2$ but $i_1 > i_2$.

An easy case analysis using Equations (5.4), (5.5), and (5.6) shows that if $m_{i_1,j_1} \succ m_{i_2,j_2}$, then the coefficient of q_{i_1,j_1} in the \mathcal{B} -representation of m_{i_2,j_2} is zero.

1. If $(i_1, j_1) = (0, 1)$, then $q_{i_1,j_1} = \alpha$, and this polynomial has a nonzero coefficient exclusively in the \mathcal{B} -representation of $m_{0,1}$.
2. The case $(i_2, j_2) = (0, 1)$ is impossible.
3. If $(i_1, j_1) = (0, 0)$, then only the fourth condition is satisfied by (i_1, j_1) , so $m_{i_1,j_1} \succ m_{i_2,j_2}$ implies $i_2 - j_2 < 1$, which in the light of (5.6) yields $j_2 \geq 2$. Consequently, by (5.5), $q_{0,0} = \beta$ has zero coefficient in the \mathcal{B} -representation of m_{i_2,j_2} .
4. If $(i_2, j_2) = (0, 0)$, then $m_{i_1,j_1} \succ m_{i_2,j_2}$ implies $i_1 - j_1 \geq 1$, so the degree of q_{i_1,j_1} is larger than the degree of m_{i_2,j_2} . Consequently, q_{i_1,j_1} has zero coefficient in the \mathcal{B} -representation of $m_{0,0}$.

5. If $i_1 - j_1 \geq 1 > i_2 - j_2$, then Equations (5.4) and (5.6) show that the three terms of m_{i_1, j_1} have higher degree than those of m_{i_2, j_2} , so in particular $\Pi(i_1 + 2, 2n - 2 - j_1)$ does not appear in the \mathcal{B} -representation of m_{i_2, j_2} .
6. If both $i_1 - j_1, i_2 - j_2 \geq 1$, then either $i_1 + 2 > i_2 + 2$, or $i_1 = i_2$ and $2n - 2 - j_1 > 2n - 2 - j_2$. In either case, by Equation (5.4), $\Pi(i_1 + 2, 2n - 2 - j_1)$ does not appear in the \mathcal{B} -representation of m_{i_2, j_2} .
7. If both $i_1 - j_1, i_2 - j_2 \leq 0$, then either $i_1 > i_2$ and $2n - j_1 = 2n - j_2$, or $2n - j_1 > 2n - j_2$. In either case, by Equation (5.4), $\Pi(i_1, 2n - j_1)$ does not appear in the \mathcal{B} -representation of m_{i_2, j_2} .

The injection $m_{i,j} \mapsto q_{i,j}$ and the linear order \succ satisfy the conditions of Corollary 5.15, therefore, by Equation (5.6),

$$\dim(\text{span}(M(\mathbb{R}^{n+2}))) = |\mathcal{B}| = (2n + 1)^2 - 2,$$

which completes the proof. □

The proof of the odd degree case is along the same lines, but it requires a considerably more complicated construction to replace the polynomials Π and the injection φ in the above proof.

Theorem 5.17. *The cone $\mathcal{P}_{2n+2}^{[a,b]}$ is not bilinear. More specifically, for every n , $\beta(\mathcal{P}_{2n+2}^{[a,b]}) = 2$.*

We have seen $\beta(\mathcal{P}_{2n+2}^{[a,b]}) \geq 2$, only the \leq direction needs proof. We first prove our claim for the case $n = 1$.

Lemma 5.18. *The cone $\mathcal{P}_4^{[a,b]}$ is not bilinear. More specifically, $\beta(\mathcal{P}_4^{[a,b]}) \leq 2$.*

Proof. Following the construction suggested by Corollary 5.9, we present a set S of 14 pairs of vectors $(\mathbf{x}, \mathbf{s}) \in C(\mathcal{P}_4^{[a,b]})$ such that the vectors $\text{vec}(\mathbf{s}\mathbf{x}^T)$ are linearly independent. Using Corollary 5.7 we can fix a and b arbitrarily; we will use $a = 1$ and $b = 6$.

For every $i = 1, \dots, 6$, let $p^{(i)}(x) = (x - 1)(x - i)^2$, $q^{(i)}(x) = (6 - x)(x - i)^2$, and define two additional polynomials $p^{(0)}(x) = (x - 1)$ and $q^{(0)}(x) = (6 - x)$. Now let S be the set consisting of the following orthogonal pairs:

- $(p^{(i)}, \mathbf{c}_{2n+2}(i)) \quad i = 1, \dots, 6,$
- $(q^{(i)}, \mathbf{c}_{2n+2}(i)) \quad i = 1, \dots, 6,$
- $(p^{(0)}, \mathbf{c}_{2n+2}(1)),$
- $(q^{(0)}, \mathbf{c}_{2n+2}(6)).$

The fact that the 14 vectors $\text{vec}(\mathbf{s}\mathbf{x}^T)$ for each $(\mathbf{x}, \mathbf{s}) \in S$ are indeed linearly independent can be verified by direct calculation (possibly with the help of a computer algebra system). \square

Proof of Theorem 5.17. If $n = 1$, then our claim is the previous lemma. From now on, let us assume $n \geq 2$. Consider the matrix valued functions $M: \mathbb{R}^{2n+2} \mapsto \mathbb{R}^{(2n+2) \times (2n+2)}$, where the entries of $M(t_1, \dots, t_n; s_1, \dots, s_n; \alpha, \beta) = (m_{i,j})_{i,j=0}^{2n+1}$ are defined as

$$M(t_1, \dots, t_n; s_1, \dots, s_n; \alpha, \beta) = \mathbf{c}_{2n+2}(t_1)\mathbf{p}^T + \mathbf{c}_{2n+2}(s_1)\mathbf{r}^T + \alpha\mathbf{c}_{2n+2}(a)\mathbf{p}_a^T + \beta\mathbf{c}_{2n+2}(b)\mathbf{p}_b^T,$$

where $\mathbf{p}, \mathbf{r}, \mathbf{p}_a, \mathbf{p}_b \in \mathcal{P}_{2n+2}^{[a,b]}$ are the coefficient vectors of the polynomials $p(x) = (x-a)\prod_{k=1}^n(x-t_k)^2$, $r(x) = (b-x)\prod_{k=1}^n(x-s_k)^2$, $p_a(x) = x-a$, and $p_b(x) = b-x$, respectively. The entries of $M = (m_{i,j})_{i,j=0}^{2n+1}$ satisfy the polynomial equation

$$\sum_{j=0}^{2n+1} m_{i,j}x^j \equiv t_1^i(x-a)\prod_{k=1}^n(x-t_k)^2 + s_1^i(b-x)\prod_{k=1}^n(x-s_k)^2 + \alpha a^i(x-a) + \beta b^i(b-x). \quad (5.7)$$

The polynomials $p(x)$, $p_a(x)$, $p_b(x)$, and $r(x)$ are nonnegative over $[a, b]$, and by Proposition 2.19,

$$\langle \mathbf{p}_a, \mathbf{c}_{2n+2}(a) \rangle = \langle \mathbf{p}_b, \mathbf{c}_{2n+2}(b) \rangle = \langle \mathbf{p}, \mathbf{c}_{2n+2}(t_1) \rangle = \langle \mathbf{r}, \mathbf{c}_{2n+2}(s_1) \rangle = 0.$$

We need to show that $\dim(\text{span}(M(\mathbb{R}^{2n+2}))) = (2n+2)^2 - 2$, then the theorem follows from Corollary 5.11. Finally, we will show this equality using the sufficient condition presented in Corollary 5.15.

Let us define the $(2n+2)$ -variate polynomials $\Pi_1(k, \ell)$ and $\Pi_2(k, \ell)$ by

$$\begin{aligned} \Pi_1(k, \ell)(t_1, \dots, t_n; s_1, \dots, s_n; \alpha, \beta) &\stackrel{\text{def}}{=} \\ &\Pi(k, \ell)(t_1, \dots, t_n) - a\Pi(k, \ell-1)(t_1, \dots, t_n), \text{ and} \\ \Pi_2(k, \ell)(t_1, \dots, t_n; s_1, \dots, s_n; \alpha, \beta) &\stackrel{\text{def}}{=} \\ &b\Pi(k, \ell-1)(s_1, \dots, s_n) - \Pi(k, \ell)(s_1, \dots, s_n), \end{aligned} \quad (5.8)$$

where $\Pi(k, \ell)$ is defined by Equation (5.3) for $0 \leq k \leq 2n+3$, $0 \leq \ell \leq 2n-2$, otherwise $\Pi(k, \ell) = 0$.

For the rest of the proof, let us assume that $a = 0$, $b = 1$; Corollary 5.7 guarantees that this is without loss of generality.

Let \mathcal{M} be the set of entries of M , and let \mathcal{B} denote the set

$$\begin{aligned} \mathcal{B} = & \{\alpha, \beta\} \cup \{\Pi(k, \ell)(t_1, \dots, t_n) \mid 3 \leq k \leq 2n+3, 0 \leq \ell \leq 2n-2, k+\ell \geq 2n+1\} \cup \\ & \cup \{\Pi(k, \ell)(s_1, \dots, s_n) \mid 3 \leq k \leq 2n+3, \ell = 2n-2\} \cup \\ & \cup \{\Pi(k, \ell)(s_1, \dots, s_n) \mid 2n \leq k \leq 2n+1, 0 \leq \ell \leq 1\} \cup \\ & \cup \{\Pi(k, \ell)(s_1, \dots, s_n) \mid 0 \leq k \leq 2n-1, 0 \leq \ell \leq 2n-2, k+\ell \leq 2n\}. \end{aligned}$$

Since $n \geq 2$, the last three sets in the union are disjoint. It follows from the definition that the set \mathcal{B} is linearly independent, because no two polynomials in \mathcal{B} share a common monomial. The coefficient of $x^{2n+1-k-\ell}$ in the polynomial $(x-a) \prod_{j=1}^n (x-t_j)^2$ is $\sum_{k=0}^2 \Pi_1(k, \ell)$. Similarly, the coefficient of $x^{2n+1-k-\ell}$ in the polynomial $(b-x) \prod_{j=1}^n (x-s_j)^2$ is $\sum_{k=0}^2 \Pi_2(k, \ell)$. From this observation and (5.7) it follows immediately that for every $0 \leq i, j \leq 2n+1$,

$$\begin{aligned} m_{i,j} = & \Pi_1(i, 2n+1-j) + \Pi_1(i+1, 2n-j) + \Pi_1(i+2, 2n-1-j) + \\ & + \Pi_2(i, 2n+1-j) + \Pi_2(i+1, 2n-j) + \Pi_2(i+2, 2n-1-j) + \\ & + m'_{i,j}, \end{aligned} \tag{5.9}$$

where $m'_{i,j}$ is defined in Equation (5.5).

We now introduce an injection $\varphi: \mathcal{B} \mapsto \mathcal{M}$ by defining its inverse (where it exists): let $m_{i,j}$ be the image of the polynomial

$$\varphi^{-1}(m_{i,j}) = q_{i,j} = \begin{cases} \alpha & i=0, j=1 \\ \Pi(i+2, 2n-1-j)(t_1, \dots, t_n) & 1 \leq j \leq \min\{i, 2n-1\} \\ \Pi(i+2, 2n-2)(s_1, \dots, s_n) & i \geq 1, j=0 \\ \beta & i=0, j=0 \\ \Pi(i, 2n+1-j)(s_1, \dots, s_n) & j \geq 3, j > \min\{i, 2n-1\} \\ \text{not defined} & \text{otherwise} \end{cases}. \tag{5.10}$$

In particular, we assign a polynomial to each entry except for $m_{0,2}$ and $m_{1,2}$, and we assign different polynomials to different entries of M , by an argument essentially identical to that in the proof of Theorem 5.16. Consequently, each polynomial in \mathcal{B} is equal to $q_{i,j}$ for at most one pair (i, j) , and Equations (5.9) and (5.10) show that (assuming $a=0$ and $b=1$) the coefficient of $q_{i,j}$ in the \mathcal{B} -representation of $m_{i,j}$ is 1 or -1 .

Let us define a linear order \succ on $\varphi(\mathcal{B})$ in the following way. Let us say that a polynomial $q_{i,j}$ is of *type* k for some $k = 1, \dots, 5$, if it is defined in the k th branch of the right-hand side of (5.10). Then, $m_{i_1, j_1} \succ m_{i_2, j_2}$ for some $(i_1, j_1) \neq (i_2, j_2)$ precisely when one of the following four conditions holds:

1. q_{i_1, j_1} is of smaller type than q_{i_2, j_2} ;
2. q_{i_1, j_1} and q_{i_2, j_2} are both of type 2, and either $i_1 > i_2$, or $i_1 = i_2$ but $j_1 < j_2$;
3. q_{i_1, j_1} and q_{i_2, j_2} are both of type 3, and $i_1 > i_2$;
4. q_{i_1, j_1} and q_{i_2, j_2} are both of type 5, and either $j_1 < j_2$, or $j_1 = j_2$ but $i_1 > i_2$.

Clearly this is indeed a linear order on $\varphi(\mathcal{B})$. An easy case analysis using (5.9), (5.5), and (5.10) shows that if $m_{i_1, j_1} \succ m_{i_2, j_2}$, then the coefficient of q_{i_1, j_1} in the \mathcal{B} -representation of m_{i_2, j_2} is zero:

1. If $(i_1, j_1) = (0, 1)$, then $q_{i_1, j_1} = \alpha$, and this polynomial has a nonzero coefficient exclusively in the \mathcal{B} -representation of $m_{0,1}$. In the remaining cases we assume $(i_1, j_1) \neq (0, 1)$.
2. The case $(i_2, j_2) = (0, 1)$ is impossible.
3. If $(i_1, j_1) = (0, 0)$, then $m_{i_1, j_1} \succ m_{i_2, j_2}$ implies $j_2 \geq 3$ (with a similar argument as in the proof of Theorem 5.16), so $q_{0,0} = \beta$ has zero coefficient in the \mathcal{B} -representation of m_{i_2, j_2} .
4. If $(i_2, j_2) = (0, 0)$, then $m_{i_1, j_1} \succ m_{i_2, j_2}$ implies $i_1 \geq 1$, so $q_{i_1, j_1} = \Pi(k, \ell)$ with some $k \geq 3$. Consequently, q_{i_1, j_1} has zero coefficient in the \mathcal{B} -representation of $m_{0,0}$. In the remaining cases we assume both (i_1, j_1) and (i_2, j_2) are different from $(0, 0)$ and $(0, 1)$.
5. If q_{i_1, j_1} and q_{i_2, j_2} are of the same type, then $i_1 > i_2$ or $j_1 < j_2$, and from Equations (5.9) and (5.10) we are done.
6. The case when q_{i_1, j_1} is of type 2 or 3 and q_{i_2, j_2} is of type 4 can be settled by comparing the degrees of q_{i_1, j_1} and m_{i_2, j_2} .
7. The only remaining case is when q_{i_1, j_1} is of type 2 and q_{i_2, j_2} is of type 3. Then $i_1 + 2 > i_2 + 2$, and hence $q_{i_1, j_1} = \Pi(i_1 + 2, 2n - 1 - j)$ has coefficient zero in the \mathcal{B} -representation of m_{i_2, j_2} .

We conclude that the injection $m_{i,j} \mapsto q_{i,j}$ and the linear order \succ satisfy the conditions of Corollary 5.15, therefore, by Equation (5.10),

$$\dim(\text{span}(M(\mathbb{R}^{2n+2}))) = |\mathcal{B}| = (2n + 2)^2 - 2,$$

which completes the proof. \square

5.1.3 Müntz polynomials

Consider a vector $\boldsymbol{\lambda} = (\lambda_0, \dots, \lambda_n)$, where $0 < \lambda_0 < \dots < \lambda_n$ are real numbers. Functions of the form $t \rightarrow \sum_{i=0}^n \alpha_i t^{\lambda_i}$ are called *Müntz polynomials of type $\boldsymbol{\lambda}$* , and they share many properties of ordinary polynomials. In particular, they form an extended Chebyshev system over $(0, \infty)$ [KS66, Chapter 1], so it is natural to ask whether the results of the previous sections generalize to the cone of Müntz polynomials nonnegative over $(0, \infty)$, denoted by

$$\mathcal{P}^{\boldsymbol{\lambda}} \stackrel{\text{def}}{=} \left\{ \mathbf{a} \mid \sum_{i=0}^n a_i t^{\lambda_i} \geq 0 \text{ for all } t \geq 0 \right\}.$$

In this section we show that the answer is at least “generically” yes. Similarly to the case of ordinary polynomials, if $t_i \geq 0$ is a root of a nonnegative Müntz polynomial $p(t)$, then it is also a root of $t \frac{d}{dt} p(t)$. This will lead to a non-trivial optimality constraint:

Proposition 5.19. $\beta(\mathcal{P}^{\boldsymbol{\lambda}}) \geq 2$.

Proof. Notice that the operator $t \frac{d}{dt}$ maps the set of Müntz polynomials of type $\boldsymbol{\lambda}$ onto itself. (Consider $\frac{d}{dt}$ to be the linear operator of formal differentiation to avoid problems when $t = 0$.) It follows from the general theory of Chebyshev systems that the nonzero extreme vectors of the dual cone of $\mathcal{P}^{\boldsymbol{\lambda}}$ are vectors of the form $\alpha(t^{\lambda_0}, \dots, t^{\lambda_n})^T$ and $\alpha(0, \dots, 0, 1)^T$, where $t \geq 0$ and $\alpha > 0$. Using the notation $\mathbf{c}(t) = (t^{\lambda_0}, \dots, t^{\lambda_n})^T$, the coefficient vector \mathbf{p} of some $p \in \mathcal{P}^{\boldsymbol{\lambda}}$ is orthogonal to $\mathbf{c}(t)$ if and only if t is a root of p . Consequently $(\mathbf{p}, \mathbf{c}(t)) \in \mathcal{C}(\mathcal{P}^{\boldsymbol{\lambda}})$ implies $p(t) = 0$; in this case either $p'(t) = 0$ or $t = 0$ must also hold, implying $t \frac{d}{dt} p(t) = 0$. The equations $p(t) = 0$ and $t \frac{d}{dt} p(t) = 0$ translate to the linearly independent bilinear complementarity relations

$$\sum_{i=0}^n p_i c_i = 0, \quad \sum_{i=0}^n \lambda_i p_i c_i = 0.$$

It is easily verified that these relations are also valid for pairs $(\mathbf{p}, \mathbf{c}) \in \mathcal{C}(\mathcal{P}^{\boldsymbol{\lambda}})$ when $\mathbf{c} = (0, \dots, 0, 1)^T$. □

Let $\Lambda_{n+1} = \{(\boldsymbol{\lambda} \in \mathbb{R}^{n+1} \mid 0 < \lambda_0 < \dots < \lambda_n)\}$ denote the space of exponent vectors for Müntz polynomials. We say that $\boldsymbol{\lambda} \in \Lambda_{n+1}$ is *generic* if the following condition holds:

*The differences $\lambda_i - \lambda_j$ are different for each pair (i, j) with $i \neq j$; furthermore, there exist indices $0 \leq i_1 < i_2 \leq n$ such that for all $j \in \{0, \dots, n\} \setminus \{i_1, i_2\}$ the value $\lambda_{i_1} + \lambda_{i_2} + \lambda_j$ can be decomposed **uniquely** as a sum of three different elements of the set $\{\lambda_0, \dots, \lambda_n\}$.*

We remark that almost all vectors in Λ_{n+1} are generic, including those with algebraically independent coordinates.

Theorem 5.20. *If λ is generic, then we have $\beta(\mathcal{P}^\lambda) = 2$ for every $n \geq 4$.*

Proof. We present a proof for $n = 2k$. The case when n is odd requires separate treatment using analogous arguments, similarly to what we have seen for ordinary polynomials over the half-line. According the previous proposition it suffices to show that $\beta(\mathcal{P}^\lambda) \leq 2$ for generic λ .

Along the lines of the previous proofs, let $M: \mathbb{R}^n \mapsto \mathbb{R}^{(2k+1) \times (2k+1)}$ be the matrix-valued function given by

$$M(t_1, \dots, t_k) = \mathbf{c}\mathbf{p}^T,$$

where $\mathbf{p} \in \mathcal{P}^\lambda$ is the coefficient vector of the Müntz polynomial p with roots t_1, \dots, t_k defined by the determinant

$$p(t) = \begin{vmatrix} t_1^{\lambda_0} & \lambda_0 t_1^{\lambda_0-1} & \dots & t_k^{\lambda_0} & \lambda_0 t_k^{\lambda_0-1} & t^{\lambda_0} \\ \vdots & & \ddots & & & \vdots \\ t_1^{\lambda_{2k}} & \lambda_{2k} t_1^{\lambda_{2k}-1} & \dots & t_k^{\lambda_{2k}} & \lambda_{2k} t_k^{\lambda_{2k}-1} & t^{\lambda_{2k}} \end{vmatrix},$$

while $\mathbf{c} = \mathbf{c}(t_1) = (t_1^{\lambda_0}, \dots, t_1^{\lambda_{2k}})$ is an extreme vector of the dual cone corresponding to the first root of \mathbf{p} .

The nonnegativity of p is a corollary of the fact that Müntz polynomials form an extended Chebyshev system over $(0, \infty)$ [KS66, Chapter 1]. Since t_1 is a root of p , we have $\langle \mathbf{p}, \mathbf{c} \rangle = p(t_1) = 0$, implying $(\mathbf{p}, \mathbf{c}) \in \mathcal{C}(\mathcal{P}^\lambda)$. Using Corollary 5.11 it only remains to show that $\dim(\text{span}(M(\mathbb{R}^k))) \geq (2k+1)^2 - 2$ holds if λ is generic.

The coefficient of t^{λ_j} in $p(t)$ is

$$\sum_{\pi \in \Pi_j} \lambda_{\pi_2} \lambda_{\pi_4} \dots \lambda_{\pi_{2k}} t_1^{\lambda_{\pi_1} + \lambda_{\pi_2} - 1} t_2^{\lambda_{\pi_3} + \lambda_{\pi_4} - 1} \dots t_k^{\lambda_{\pi_{2k-1}} + \lambda_{\pi_{2k}} - 1},$$

where Π_j denotes the family of all permutations of the set $\{0, \dots, 2k\} \setminus \{j\}$. The (i, j) entry of the matrix $M(t_1, \dots, t_k)$ can then be expressed as

$$m_{i,j} = \sum_{\pi \in \Pi_j} \lambda_{\pi_2} \lambda_{\pi_4} \dots \lambda_{\pi_{2k}} t_1^{\lambda_i + \lambda_{\pi_1} + \lambda_{\pi_2} - 1} t_2^{\lambda_{\pi_3} + \lambda_{\pi_4} - 1} \dots t_k^{\lambda_{\pi_{2k-1}} + \lambda_{\pi_{2k}} - 1}.$$

Since $\lambda > 0$, all monomials in the above formula have a nonzero coefficient.

The sum of the exponents in each monomial of entry $m_{i,j}$ is $\lambda_i - \lambda_j + \sum_{\ell=0}^{2k} \lambda_\ell - k$. Let $\boldsymbol{\lambda}$ be generic. Since the pairwise differences $\lambda_i - \lambda_j$ ($i \neq j$) are different, the non-diagonal entries $m_{i,j}$, when viewed as Müntz polynomials of variables t_1, \dots, t_k , are all linearly independent, as they do not share a common monomial. Furthermore, since $\boldsymbol{\lambda}$ is generic, monomials containing the term $t_1^{\lambda_{i_1} + \lambda_{i_2} + \lambda_j}$ do not appear in diagonal entries $m_{i,i}$, unless $i \in \{i_1, i_2, j\}$. Hence the diagonal entries $m_{j,j}$ satisfying $j \neq i_1$ and $j \neq i_2$, and the off-diagonal entries are all linearly independent Müntz polynomials, and $\dim(\text{span}(M(\mathbb{R}^k))) \geq (2k+1)^2 - 2$; that is, $\beta(\mathcal{P}^\lambda) \leq 2$ in the generic case. \square

5.1.4 Cones linearly isomorphic to nonnegative polynomials over an interval

As we have seen in Proposition 5.6, applying a nonsingular linear transformation A to a cone \mathcal{K} preserves its bilinearity rank. In addition, any change of basis for the cone of polynomials will result in a cone linearly isomorphic to it. Thus, for instance, the set of vectors of coefficients of nonnegative polynomials expressed in any orthogonal polynomial basis (e.g., Laguerre, Legendre, Chebyshev, etc.), or the Bernstein polynomial basis are linearly isomorphic to the cone of nonnegative polynomials in the standard basis. This fact is useful in numerical computations since the standard basis is numerically unstable and we may need to work with a more stable basis.

We have already stated in Corollary 5.7 that for all $a < b$ and n , the cones $\mathcal{P}_{n+1}^{[a,b]}$ and $\mathcal{P}_{n+1}^{[0,1]}$ are linearly isomorphic. More generally:

Lemma 5.21. *Let f be a function with domain $\Delta \subseteq \mathbb{R}$ and range $\Omega \subseteq \mathbb{R}$; also suppose that the set of functions $\{1, f, f^2, \dots, f^n\}$ is linearly independent. Then the cone*

$$\mathcal{P}^f = \left\{ \mathbf{a} = (a_0, \dots, a_n) \left| \sum_{j=0}^n a_j f^j(t) \geq 0 \text{ for all } t \in \Delta \right. \right\}$$

is linearly isomorphic to the cone of ordinary polynomials nonnegative over Ω .

From this observation, and using change of basis as needed, we can prove linear isomorphism of a number of cones of nonnegative functions over well-known finite-dimensional bases with \mathcal{P}_{2n+1} or $\mathcal{P}_{n+1}^{[0,1]}$. Below we present a partial list. Most of the techniques used below are quite simple, and they are used by Karlin and Studden [KS66] and Nesterov [Nes97] for other purposes.

Rational functions

The basis $\{t^{-m}, t^{-m+1}, \dots, t^{n-1}, t^n\}$ for nonnegative even integers n and m spans the set of rational functions with a degree n numerator and denominator t^m . Since $\sum_{i=-m}^n p_i t^i = t^{-m} \sum_{i=0}^{n+m} p_i t^i$,

the cone of rational functions with numerator of degree n and denominator t^m nonnegative over $\Delta = \mathbb{R} \setminus \{0\}$ is linearly isomorphic to the cone of nonnegative polynomials of degree $n + m$, and therefore its bilinearity rank is 4.

Nonnegative polynomials over $[0, \infty)$

Consider the basis $B = \{t^n, t^{n-1}(1-t), \dots, t(1-t)^{n-1}, (1-t)^n\}$ of polynomials of degree n . Clearly the cone

$$\left\{ (p_0, \dots, p_n) \left| \sum_{i=0}^n p_i t^i (1-t)^{n-i} \geq 0 \quad \forall t \in [0, 1] \right. \right\},$$

which consists of coefficient vectors of polynomials nonnegative over $[0, 1]$, expressed in basis B , is linearly isomorphic to $\mathcal{P}_{n+1}^{[0,1]}$. On the other hand we have:

Proposition 5.22 ([Nes97]). *A polynomial $p_0(1-t)^n + p_1t(1-t)^{n-1} + \dots + p_nt^n$ is nonnegative over $[0, 1]$ if and only if the polynomial $p_0 + p_1t + \dots + p_nt^n$ is nonnegative over $[0, \infty)$.*

This follows from

$$\sum_k p_k t^k (1-t)^{n-k} = (1-t)^n \sum_k p_k \left(\frac{t}{1-t} \right)^k$$

and the fact that $[0, 1]$ is mapped to $[0, \infty)$ under $f(t) = \frac{t}{1-t}$. We thus get the following result.

Corollary 5.23. *The cone $\mathcal{P}_{n+1}^{[0,\infty)}$ is linearly isomorphic to $\mathcal{P}_{n+1}^{[0,1]}$. Therefore $\beta(\mathcal{P}_{n+1}^{[0,\infty)}) = 2$ for every positive integer n .*

Cosine polynomials

Consider the cone

$$\mathcal{P}_{n+1}^{\cos} \stackrel{\text{def}}{=} \left\{ \mathbf{c} \in \mathbb{R}^{n+1} \left| \sum_{k=0}^n c_k \cos(kt) \geq 0 \text{ for all } t \in \mathbb{R} \right. \right\}.$$

To relate this cone to the cones we have discussed before, first observe that $\cos(kt)$ can be expressed as an ordinary polynomial of degree k of $\cos(t)$ using the well-known Chebyshev polynomials of the first kind [MH03]. This follows immediately from applying the binomial theorem to the identity $(\cos(t) + i \sin(t))^k = \cos(kt) + i \sin(kt)$:

$$\cos(kt) = \sum_{j=0}^{\lfloor k/2 \rfloor} \binom{k}{2j} \cos^{k-2j}(t) (1 - \cos^2(t))^j \tag{5.11a}$$

$$\sin(kt) = \sin(t) \sum_{j=0}^{\lfloor k/2 \rfloor} \binom{k}{2j+1} \cos^{k-2j-1}(t) (1 - \cos^2(t))^j \quad (5.11b)$$

From (5.11a) we see that $\cos(kt)$ is a polynomial of $\cos(t)$ of degree k . (This polynomial is known as the k -th Chebyshev polynomial of the first kind; see also Definition 5.28.) Thus, every vector \mathbf{c} representing the cosine polynomial $c(t) = \sum_{k=0}^n c_k \cos(kt)$ is mapped to a vector \mathbf{p} representing the ordinary polynomial $p(s) = \sum_{k=0}^n p_k s^k$ through the identity $\sum_{k=0}^n c_k \cos(kt) = p(\cos(t))$. Furthermore, this correspondence between \mathbf{c} and \mathbf{p} is one-to-one and onto, since for each k the function $\cos(kt)$ is a polynomial of degree k in $\cos(t)$ the matrix mapping \mathbf{c} to \mathbf{p} is lower triangular with nonzero diagonal entries. Now $c(t) = p(\cos(t)) \geq 0$ for all t if and only if $p(s) \geq 0$ for all $s \in [-1, 1]$. Recalling that $\mathcal{P}_{n+1}^{[-1,1]}$ is linearly isomorphic to $\mathcal{P}_{n+1}^{[0,1]}$, we have the following result.

Corollary 5.24. *The cone \mathcal{P}_{n+1}^{\cos} is linearly isomorphic to $\mathcal{P}_{n+1}^{[0,1]}$. Therefore $\beta(\mathcal{P}_{n+1}^{\cos}) = 2$ for every $n \in \mathbb{N}$.*

Trigonometric polynomials

Consider the cone

$$\begin{aligned} \mathcal{P}_{2n+1}^{\text{trig}} &= \left\{ \mathbf{r} \in \mathbb{R}^{2n+1} \mid r_0 + \sum_{k=1}^n (r_{2k-1} \cos(kt) + r_{2k} \sin(kt)) \geq 0 \text{ for all } t \in \mathbb{R} \right\} \\ &= \left\{ \mathbf{r} \in \mathbb{R}^{2n+1} \mid r_0 + \sum_{k=1}^n (r_{2k-1} \cos(kt) + r_{2k} \sin(kt)) \geq 0 \text{ for all } t \in (-\pi, \pi) \right\} \end{aligned}$$

To transform a trigonometric polynomial $r(t) = r_0 + \sum_{k=1}^n (r_{2k-1} \cos(kt) + r_{2k} \sin(kt))$ into one of the classes of polynomials already discussed we make a change of variables $t = 2 \arctan(s)$. With this transformation we have

$$\begin{aligned} \sin(t) &= \frac{2s}{1+s^2} \\ \cos(t) &= \frac{1-s^2}{1+s^2} \end{aligned}$$

Using (5.11) we can write

$$r(t) = T_n\left(\frac{1-s^2}{1+s^2}\right) + \frac{2s}{1+s^2} U_{n-1}\left(\frac{1-s^2}{1+s^2}\right) \quad (5.12)$$

where T_n and U_{n-1} are ordinary polynomials of degree n , and $n-1$, respectively (the Chebyshev polynomials of the first and second kind [MH03]); T_n is obtained from (5.11a) and U_{n-1} is obtained

from (5.11b). Multiplying by $(1 + s^2)^n$ we see that

$$r(t) = (1 + s^2)^{-n} p(s)$$

for some ordinary polynomial p . Substituting (5.11a) and (5.11b), the polynomial p can be expressed in the following basis:

$$\left\{ (1 + s^2)^n, (1 + s^2)^{n-1}(1 - s^2), \dots, (1 - s^2)^n \right\} \cup \left\{ s(1 + s^2)^{n-1}, s(1 - s^2)^{n-2}(1 - s^2), \dots, s(1 - s^2)^{n-1} \right\}$$

It is straightforward to see that this is indeed a basis. We need to simply observe that those terms that are not multiplied by s form a basis of polynomials with even degree terms, and those that involve s form a basis of polynomials with odd degree terms. Therefore, the correspondence between vector of coefficients \mathbf{r} of the trigonometric polynomial $r(t)$ and the vector of coefficients \mathbf{p} of the ordinary polynomial $p(s)$ in the above basis is one-to-one and onto. Furthermore, since the function $\tan(t/2)$ maps $(-\pi, \pi)$ to \mathbb{R} , it follows that trigonometric and ordinary polynomials of the same degree are linearly isomorphic.

Corollary 5.25. *The cone $\mathcal{P}_{2n+1}^{trig}$ is linearly isomorphic to \mathcal{P}_{n+1} . Therefore, $\beta(\mathcal{P}_{2n+1}^{trig}) = 4$ for every $n \in \mathbb{N}$.*

Exponential polynomials

One could ask if the results of trigonometric polynomials extend to hyperbolic functions \sinh and \cosh or to exponential functions. The situation is actually simpler here. Consider the cone

$$\mathcal{P}^{\text{exp}} \stackrel{\text{def}}{=} \left\{ \mathbf{e} \mid \sum_{k=-m}^n e_k \exp(kt) \geq 0 \text{ for all } t \geq 0 \right\}.$$

First, there is no loss of generality if we assume $m = 0$ since every such polynomial can be multiplied by $\exp(mt)$. Now clearly $e(t) = e_0 + e_1 \exp(t) + \dots + e_n \exp(nt) \geq 0$ for all $t \in \mathbb{R}$ if and only if the ordinary polynomial $e_0 + e_1 s + \dots + e_n s^n$ is nonnegative over $[0, \infty]$. Recalling that $\mathcal{P}^{[0, \infty]}$ is linearly isomorphic to $\mathcal{P}^{[0, 1]}$ we have shown the following:

Corollary 5.26. *The cone \mathcal{P}^{exp} is linearly isomorphic to $\mathcal{P}^{[0, 1]}$. Therefore, $\beta(\mathcal{P}^{\text{exp}}) = 2$.*

5.2 A barrier method for moment cone constraints

In this section we return to barrier methods tailored for nonnegative polynomials, which may be used to circumvent the use of semidefinite programming in solving polynomial optimization problems. We concentrate on the univariate case, generalizing the results of this section to the multivariate setting is subject of future research.

Consider a conic optimization problem involving K moment cones in the following form:

$$\begin{aligned} & \text{minimize} && \sum_i \langle \mathbf{c}_i, \mathbf{x}_i \rangle \\ & \text{subject to} && \sum_i \mathbf{A}_i \mathbf{x}_i = \mathbf{b} \\ & && \mathbf{x}_i \in \mathcal{M}_{2n+1} \quad i = 1, \dots, K \end{aligned}$$

Using the semidefinite representation of \mathcal{M}_{2n+1} (Proposition 2.18), the equivalent semidefinite program is obtained by replacing the conic constraints with

$$H(\mathbf{x}_i) \in \mathbb{S}_+^{n+1} \quad i = 1, \dots, K,$$

where $H(\mathbf{x})$ is the Hankel matrix defined in (2.6), that is, $H(\mathbf{x}) = (x_{i+j})_{i,j=0,\dots,n}$.

Path-following interior-point methods solve the above semidefinite programming problem by following (approximately) the curve consisting of the (unique) minimizers of the constrained convex problems

$$\begin{aligned} & \text{minimize} && \sum_i \langle \mathbf{c}_i, \mathbf{x}_i \rangle + \mu \ln(\det(H(\mathbf{x}_i))) \\ & \text{subject to} && \sum_i \mathbf{A}_i \mathbf{x}_i = \mathbf{b} \end{aligned}$$

in the direction $\mu \searrow 0$. In each iteration of the (short-step) path-following algorithms the vectors \mathbf{x}_i are updated by reducing μ by some factor and taking one Newton step from the current approximate solution towards the optimal solution of the problem obtained using the reduced μ . The bottleneck of the Newton step is the solution of an unstructured system of a linear equations, which takes $\mathcal{O}(n^3)$ time. However, setting up the equations requires an evaluation of the objective function of the above problem, its gradient, and its Hessian—which takes $\mathcal{O}(n^6)$ time without using the special structure of $H(\mathbf{x})$. Hence, the difficulty in accelerating the path-following interior-point methods lies

in computing the gradient and Hessian of the objective function fast, preferably in subcubic time, which reduces to computing the gradient $\nabla_{\mathbf{x}}f$ and the Hessian $\nabla_{\mathbf{x}}^2f$ of the function $f = \ln(\det(H(\cdot)))$ for arguments \mathbf{x} that yield a positive definite Hankel matrix $H(\mathbf{x})$.

Other dual methods, such as long-step barrier methods also require the computation of the gradient and the Hessian of f as one of the key steps of each iteration of the algorithm. By carrying out these computations in subcubic time, the complexity of every iteration of each of these algorithms also reduces to $O(n^3)$.

5.2.1 Fast computation of the gradient and the Hessian

It is well-known (and easily established) that the gradient of the function $\mathbf{X} \rightarrow \ln(\det(\mathbf{X}))$ (where $\mathbf{X} \in \mathbb{R}^{n \times n}$ is a matrix with positive determinant) is given by the formula

$$\nabla_{\mathbf{X}} \ln(\det(\mathbf{X})) = (\mathbf{X}^T)^{-1}.$$

We need to apply this formula to matrices of the form $\mathbf{X} = H(\mathbf{x}) = \sum_{i=0}^{2n} \mathbf{E}_i x_i$, where $\mathbf{E}_i \in \{0, 1\}^{(n+1) \times (n+1)}$ ($i = 0, \dots, 2n$) is the matrix having a 1 at those positions (j, k) which satisfy $j + k = i$.

Let $\mathbf{X} = H(\mathbf{x})$, where $\mathbf{x} = (x_0, \dots, x_{2n})^T$. Using the chain rule we obtain

$$\frac{\partial}{\partial x_i} \ln(\det(H(\mathbf{x}))) = \left\langle \nabla_{\mathbf{X}} \ln(\det(\mathbf{X})), \frac{\partial}{\partial x_i} H(\mathbf{x}) \right\rangle = \langle H(\mathbf{x})^{-1}, \mathbf{E}_i \rangle, \quad i = 0, \dots, 2n.$$

That is, the i th entry of $\nabla_{\mathbf{x}}H(\mathbf{x})$ is the sum of the i th reverse diagonal entries of the inverse matrix of $H(\mathbf{x})$. Special algorithms are available to compute the inverse of a Hankel matrix in $\mathcal{O}(n^2)$ time and space, or even faster, see below; consequently $\nabla_{\mathbf{x}}H(\mathbf{x})$ can be computed in $\mathcal{O}(n^2)$ time.

To obtain the Hessian, we differentiate $\nabla_{\mathbf{x}}H(\mathbf{x})$ with respect to \mathbf{x} . The identity for the derivative of the inverse in differential form can be written as follows:

$$d\mathbf{X}^{-1} = -\mathbf{X}^{-1}(d\mathbf{X})\mathbf{X}^{-1},$$

which yields that

$$d(H(\mathbf{x})^{-1}) = -H(\mathbf{x})^{-1} \left(d \left(\sum_j \mathbf{E}_j x_j \right) \right) H(\mathbf{x})^{-1} = - \sum_j (H(\mathbf{x})^{-1} \mathbf{E}_j H(\mathbf{x})^{-1}) dx_j,$$

that is, for every pair (i, j) ,

$$\frac{\partial^2}{\partial x_i \partial x_j} \ln(\det(H(\mathbf{x}))) = -\langle H(\mathbf{x})^{-1} \mathbf{E}_j H(\mathbf{x})^{-1}, \mathbf{E}_i \rangle, \quad i, j = 0, \dots, 2n.$$

Let z_{ij} denote the (i, j) -th entry of $H(\mathbf{x})^{-1}$. Then the matrix $H(\mathbf{x})^{-1} \mathbf{E}_j H(\mathbf{x})^{-1}$ can be written as

$$H(\mathbf{x})^{-1} \mathbf{E}_j H(\mathbf{x})^{-1} = \left(\sum_{k+l=j} z_{ak} z_{lb} \right)_{a,b=0,\dots,2n},$$

which in turn yields

$$-\frac{\partial^2}{\partial x_i \partial x_j} \ln(\det(H(\mathbf{x}))) = \langle H(\mathbf{x})^{-1} \mathbf{E}_j H(\mathbf{x})^{-1}, \mathbf{E}_i \rangle = \sum_{\substack{a+b=i \\ k+l=j}} z_{ak} z_{lb} = \sum_{\substack{a+b=i \\ k+l=j}} z_{ak} z_{bl},$$

where the last equality holds because $H(\mathbf{x})^{-1}$ is symmetric.

The last summation shows that the negative of the Hessian is the convolution of $H(\mathbf{x})^{-1}$ by itself; equivalently, $-\nabla_{\mathbf{x}}^2 \ln(\det(H(\mathbf{x})))$ is the coefficient matrix of the square of the bivariate polynomial whose coefficient matrix is $H(\mathbf{x})^{-1}$. If $H(\mathbf{x})^{-1}$ is already computed, then this convolution can be computed by a single bivariate polynomial multiplication.

The complete algorithm is summarized in the following theorem and its proof.

Theorem 5.27. *The gradient and the Hessian of the function $\mathbf{x} \rightarrow \ln(\det(H(\mathbf{x})))$ can be computed in $\mathcal{O}(n^2 \log n)$ time.*

Proof. The inverse of the Hankel matrix $H(\mathbf{x})$ can be computed in $\mathcal{O}(n^2)$ time by practical special purpose algorithms such as the ones in [HJ90] and [Pan01, Chapter 5]. (In fact, the inverse of a Hankel matrix can even be computed in $\mathcal{O}(n \log^2 n)$ time, see the same sources, but the algorithms that achieve subquadratic complexity, sometimes referred to as “superfast” algorithms, are only asymptotically fast, and are not practical, unless n is extremely large.)

Based on the above calculations, we compute the i th component of the gradient by adding the entries of $H(\mathbf{x})^{-1}$ along the i th reverse diagonal ($i = 0, \dots, 2n$). To find the Hessian, consider the polynomial $p(x, y) = \sum_{i,j=0}^n z_{ij} x^i y^j$, where z_{ij} is the (i, j) -th entry of $H(\mathbf{x})^{-1}$. Compute the coefficients q_{ij} of the polynomial

$$q(x, y) = -p^2(x, y) = \sum_{i,j=0}^{2n} q_{ij} x^i y^j.$$

The matrix $(q_{ij})_{i,j=0,\dots,2n}$ is the Hessian.

The bottleneck of the algorithm is the computation of $p^2(x, y)$, which can be carried out in $\mathcal{O}(n^2 \log n)$ time using two-dimensional fast Fourier transform. \square

5.2.2 Alternative bases

Representing polynomials in the standard basis frequently results in serious numerical problems when the degree of the polynomials is high enough. In practice, current SDP solvers that use IEEE double precision arithmetic reliably handle positive polynomial constraints up to degree 8-12. From the dual point of view, Hankel matrices are known to have large condition numbers [Tyr94, Bec97], and computing the inverse of a positive definite Hankel matrix in every iteration of the algorithm is problematic even if we use special purpose algorithms that exploit the Hankel structure. Consequently moment conic constraints of high dimension cannot be handled easily.

To alleviate the numerical difficulties we can represent the positive polynomials in some other basis known to be more suitable for numerical computations, for example, in some orthogonal basis. The proof of Theorem 5.27 carries over to any other basis as long as the following hold:

- the square of a polynomial, expressed in the new basis; either directly, or via a basis transformation from the new basis to the standard monomial basis, can be computed in $\mathcal{O}(n^2 \log n)$ time;
- the computations involving Hankel matrices can be replaced by analogous computations involving some other family of structured matrices, for which matrix inversion can be carried out in $\mathcal{O}(n^2 \log n)$ time.

One such basis is the Chebyshev polynomial basis.

Definition 5.28. Let T_n be the n th Chebyshev polynomial of the first kind, defined by the recursion

$$T_0(x) = 1, \quad T_1(x) = x, \quad T_n(x) = 2xT_{n-1}(x) - T_{n-2}(x) \quad (n \geq 2).$$

For this orthogonal basis we define the cone of *nonnegative Chebyshev polynomials* as

$$\mathcal{P}_{2n+1}^{\text{Ch}} = \left\{ (p_0, \dots, p_{2n})^T \mid \sum_{i=0}^{2n} p_i T_i(x) \geq 0 \quad \forall x \in \mathbb{R} \right\},$$

and the cone of *Chebyshev moments* as

$$\mathcal{M}_{2n+1}^{\text{Ch}} = \text{cl cone} \left(\{ (T_0(x), \dots, T_{2n}(x))^T \mid x \in \mathbb{R} \} \right).$$

Clearly, these cones are the duals of each other, and they are simply linear transformations of \mathcal{P}_{2n+1} and \mathcal{M}_{2n+1} , respectively, hence they are semidefinite representable. We have the following theorem, analogously to the well-known semidefinite representation of positive polynomials and moments.

Theorem 5.29. *Let $\Lambda: \mathbb{R}^{2n+1} \rightarrow \mathbb{R}^{(n+1) \times (n+1)}$ be the linear operator satisfying $\Lambda(\mathbf{c}) = (\lambda_{ij})_{i,j=0,\dots,n}$ with*

$$\lambda_{ij} = \frac{c_{i+j} + c_{|i-j|}}{2} \quad i, j = 0, \dots, n,$$

and let Λ^ be its adjoint. Then for every $\mathbf{p} \in \mathbb{R}^{2n+1}$, $\mathbf{p} \in \mathcal{P}_{2n+1}^{\text{Ch}}$ if and only if there exists a positive semidefinite matrix \mathbf{X} such that $\Lambda^*(\mathbf{X}) = \mathbf{p}$. Furthermore, for every $\mathbf{c} \in \mathbb{R}^{2n+1}$, $\mathbf{c} \in \mathcal{M}_{2n+1}^{\text{Ch}}$ if and only if the matrix $\Lambda(\mathbf{c})$ is positive semidefinite.*

Proof. It is possible to derive this result from the semidefinite representation of positive polynomials represented in the standard basis (recall Proposition 2.3). Alternatively, since nonnegative polynomials are the same as sum-of-squares polynomials (regardless of the basis they are represented in), the claim is a special case of Theorem 4.5. To complete the proof we need construct the operator Λ as in (4.9). Using the identity $T_i T_j = \frac{1}{2}(T_{i+j} + T_{|i-j|})$, see [MH03, Section 2.4], we obtain immediately that the Λ defined in (4.9) is the one defined in the claim. The characterization of the dual cone comes from Theorem 4.6. \square

Notice that the matrix $\Lambda(\mathbf{c})$ in Theorem 5.29 is a *Toeplitz+Hankel matrix*, that is, it can be written as a sum of a Toeplitz matrix and a Hankel matrix. The inverse of a Toeplitz+Hankel matrix can be computed in $O(n^2)$ time [Pan01]. Furthermore, extensions of the fast Fourier transform to Chebyshev polynomials are also known [MH03, Chapter 4]. Consequently, the algorithm given in the proof of Theorem 5.27 carries over to positive Chebyshev polynomials and Chebyshev moments, and the running times are the same.

5.3 Conclusion

Our goal was to see if there is a way to solve optimization problems involving conic constraints of the form $\mathbf{p} \in \mathcal{P}_{n+1}^{[a,b]}$ or $\mathbf{c} \in \mathcal{M}_{n+1}^{[a,b]}$ without formulating them as semidefinite programs. This question led us to consider first the simplest form of complementarity relations, bilinear equations, for nonnegative polynomials and moment vectors. For the first orthant, the second-order cone, and the semidefinite cone bilinear complementary relations suffice to represent complementarity, and this leads to efficient

homotopy methods for the corresponding conic programs. However, we have found that bilinear relations alone are not sufficient to represent complementarity between nonnegative polynomials and moment vectors; in fact, no more than two linearly independent bilinear relations hold for pairs from the complementarity set $\mathcal{C}(\mathcal{P}_{n+1}^{[a,b]})$. We have found similar theorems for a number of related cones.

In the second half of the chapter we turned towards barrier methods. The universal barrier function (and its derivatives) appear to be very difficult to evaluate for nonnegative polynomials and moments. On the other hand, we have found that the logarithmic barrier function of the semidefinite cone, along with its gradient and Hessian, can be evaluated for Hankel matrices and Toeplitz+Hankel matrices very fast. Hence, moment cone constraints can be handled efficiently by simulating SDP interior point iterations and projections to the space of Hankel matrices in subcubic time and space, without carrying out the lifting to the space of symmetric matrices.

Chapter 6

Conclusions and Open Problems

We have seen that the semidefinite and second-order cone representation of nonnegative polynomials provide practical models (that are also theoretically efficient) for several shape-constrained spline estimation problems. In practice, these models are applicable to very large scale univariate estimation problems with bound constraints, monotonicity and convexity constraints, interpolation and periodicity constraints on the estimator, even combined with cross-validation of parameters. We demonstrated that polyhedral approximations, while popular in the statistics literature, are not necessary for such statistical estimation problems, as they can be solved directly and equally efficiently.

Polyhedral approximations to these nonlinear models have been proposed by several authors in the literature, sometimes without justification that the approximate models indeed provide good approximation. We examined this problem, and gave a necessary and sufficient condition for a cone of piecewise polynomial splines with nonnegative coordinates in a nonnegative basis to be dense in the cone of nonnegative continuous functions.

In the multivariate setting such approximations are unavoidable on complexity grounds. We considered both polyhedral and spectrahedral (semidefinite) approximations derived from weighted-sum-of-squares representation of nonnegative polynomials. With a generalization of the univariate density result, theoretical justification was provided for the approximations suggested.

Theoretically, splines of any degree and order of differentiability can be used regardless of the objective of the estimation (least-squares, maximum likelihood, or maximum penalized likelihood), but the lack of off-the-shelf modeling tools and solvers simultaneously supporting arbitrary convex objective functions and semidefinite constraints makes some of the models somewhat less practical

than others. This is especially true for multivariate problems.

Even when ready-made solvers are available, algorithms designed specifically for the estimation problems could greatly increase the practicality of the proposed models, especially in the multivariate case. The thesis contains some results in this direction, but several questions remain open. A better understanding of low-degree nonnegative polynomials over low-dimensional polyhedra could lead to better approximations. Experimenting with decomposition schemes different from the one proposed in the thesis could also lead to even more efficient solutions of large-scale multivariate models.

We have seen that some constraints on the shape of the estimator cannot be translated to nonnegativity. Laying the groundwork for future computational research in this area, we presented a semidefinite characterization of sum-of-squares and weighted-sum-of-squares cones that is considerably more general than those in the literature. We hope that this result will also have an audience outside the area of function estimation.

Appendix A

Background material

This appendix is a summary some of the fundamentals of conic optimization, positive semidefiniteness, and semidefinite programming, which are used throughout the thesis. A large number introductory texts and survey papers are available in each of these topics, see for example [HJ85], [WSV00], and [Ren01].

A.1 Conic optimization

Definition A.1 (cone, convex cone). A non-empty subset \mathcal{K} of a real linear space V is called a *cone* if for every $\lambda \geq 0$ and $\mathbf{x} \in V$, $\mathbf{x} \in \mathcal{K}$ implies $\lambda\mathbf{x} \in \mathcal{K}$; \mathcal{K} is a *convex cone*, if it is also convex, for which it is sufficient that $\mathbf{x} + \mathbf{y} \in \mathcal{K}$ whenever $\mathbf{x} \in \mathcal{K}$ and $\mathbf{y} \in \mathcal{K}$.

Definition A.2 (vector order, positive cone). A *vector order* of a real linear space V is a partial order \succcurlyeq such that $\mathbf{x} + \mathbf{z} \succcurlyeq \mathbf{y} + \mathbf{z}$ and $\lambda\mathbf{x} \succcurlyeq \lambda\mathbf{y}$ for every $\lambda \geq 0$ and $\mathbf{x}, \mathbf{y}, \mathbf{z} \in V$ satisfying $\mathbf{x} \succcurlyeq \mathbf{y}$. If \succcurlyeq is a vector order, then the set $\{\mathbf{x} \mid \mathbf{x} \succcurlyeq 0\}$ is a convex cone, and it is called the *positive cone* of \succcurlyeq . Conversely, if \mathcal{K} is a convex cone, then we can define a relation $\succcurlyeq_{\mathcal{K}}$ by setting $\mathbf{x} \succcurlyeq_{\mathcal{K}} \mathbf{y}$ if and only if $\mathbf{x} - \mathbf{y} \in \mathcal{K}$. This is a partial order, and we say that it *corresponds to* the convex cone \mathcal{K} . The relation $\mathbf{x} \succcurlyeq_{\mathcal{K}} \mathbf{y}$ can also be written as $\mathbf{y} \preccurlyeq_{\mathcal{K}} \mathbf{x}$.

The correspondence in the above definition is one-to-one.

Proposition A.3 ([KN63, §1.2.4]). *Every vector order is the partial order corresponding to its positive cone. Conversely, every convex cone is the positive cone of its corresponding partial order.*

The partial order $\succcurlyeq_{\mathcal{K}}$ is not particularly meaningful unless $\mathbf{0} \neq \mathbf{x} \succcurlyeq 0$ prohibits $-\mathbf{x} \succcurlyeq 0$. Cones with this property are called *pointed*.

Consider now a space V equipped with an inner product $\langle \cdot, \cdot \rangle$. For the purposes of this thesis it is sufficient to consider cones in finite-dimensional spaces, hence, we will assume that $V = \mathbb{R}^n$ for some $n \geq 1$. Consequently, the inner product is of the form $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^T \mathbf{Q} \mathbf{y}$ for some positive definite matrix \mathbf{Q} . The basic object of conic optimization (also called conic programming) is the *conic optimization problem in standard form*:

$$\begin{aligned} & \text{minimize} && \langle \mathbf{c}, \mathbf{x} \rangle \\ & \text{subject to} && \mathbf{A} \mathbf{x} = \mathbf{b}, \\ & && \mathbf{x} \succ_{\mathcal{K}} \mathbf{0} \end{aligned} \tag{A.1}$$

where the matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and the vectors $\mathbf{b} \in \mathbb{R}^m$ and $\mathbf{c} \in \mathbb{R}^n$ are given. If we choose $\mathcal{K} = \mathbb{R}_+^n$, the first orthant, we obtain the familiar linear programming problem. The objective “minimize” must be understood in the “inf” sense, since (unlike in linear programming) the minimum may not exist even if the set $\{\mathbf{x} \mid \mathbf{A} \mathbf{x} = \mathbf{b}\} \cap \mathcal{K}$ is non-empty, and \mathcal{K} is closed.

It is usually assumed that \mathcal{K} is *closed*. Furthermore, it is often convenient to assume that it is *full-dimensional*, that is $\text{span}(\mathcal{K}) = V = \mathbb{R}^n$. Note that this is not an intrinsic property of the cone itself, but it depends on its representation, in other words, on the space it is embedded in. Closed, convex, full-dimensional cones are called *proper cones*.

The *dual* of the optimization problem (A.1) can also be cast as a conic optimization problem involving the *dual cone* of \mathcal{K} , denoted by \mathcal{K}^* :

$$\mathcal{K}^* \stackrel{\text{def}}{=} \{\mathbf{y} \mid \langle \mathbf{x}, \mathbf{y} \rangle \geq 0 \quad \forall \mathbf{x} \in \mathcal{K}\}.$$

The name “dual cone” is justified by the facts that \mathcal{K}^* is indeed a cone, and that if \mathcal{K} is a proper cone, then $(\mathcal{K}^*)^* = \mathcal{K}$. (Note that the dual cone also depends on the space \mathcal{K} is embedded in.) The dual of (A.1) is the following optimization problem:

$$\begin{aligned} & \text{maximize} && \langle \mathbf{b}, \mathbf{y} \rangle \\ & \text{subject to} && \mathbf{A}^T \mathbf{y} \preceq_{\mathcal{K}^*} \mathbf{c} \end{aligned} \tag{A.2}$$

Equivalently, we can write the dual problem in the following form:

$$\begin{aligned} & \text{maximize} && \langle \mathbf{b}, \mathbf{y} \rangle \\ & \text{subject to} && \mathbf{c} - \mathbf{A}^T \mathbf{y} = \mathbf{s} \\ & && \mathbf{s} \succ_{\mathcal{K}^*} \mathbf{0} \end{aligned} \tag{A.2'}$$

As in linear programming, the original problem (A.1) is commonly referred to as the *primal* problem.

Similarly to the duality theory of linear programming, we have the following simple fact.

Proposition A.4. *If \mathbf{x} is a feasible solution to (A.1), and (\mathbf{y}, \mathbf{s}) is a feasible solution to (A.2'), then*

$$\langle \mathbf{c}, \mathbf{x} \rangle - \langle \mathbf{b}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{s} \rangle \geq 0.$$

As a corollary, we have a sufficient condition for optimality.

Proposition A.5. *If \mathbf{x} is a feasible solution to (A.1), (\mathbf{y}, \mathbf{s}) is a feasible solution to (A.2'), and the complementarity condition $\langle \mathbf{x}, \mathbf{s} \rangle = 0$ holds, then \mathbf{x} and (\mathbf{y}, \mathbf{s}) are optimal solutions to the primal problem (A.1) and dual problem (A.2'), respectively.*

A.2 Positive semidefinite matrices

We only consider real symmetric positive semidefinite matrices. They are characterized by several key properties; the ones used in the proofs in this thesis are summarized in the following proposition.

Proposition A.6 ([HJ85, Chapter 7]). *Let \mathbf{A} be a real symmetric $n \times n$ matrix. Then the following properties are equivalent.*

1. For every $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{x}^T \mathbf{A} \mathbf{x} \geq 0$.
2. All eigenvalues of \mathbf{A} are nonnegative. (Recall that all eigenvalues of a real symmetric matrix are real numbers.)
3. There exists an $n \times m$ matrix \mathbf{B} such that $\mathbf{A} = \mathbf{B} \mathbf{B}^T$.
4. There exists an $n \times n$ lower triangular matrix \mathbf{L} such that $\mathbf{A} = \mathbf{L} \mathbf{L}^T$.
5. The matrix \mathbf{A} is a nonnegative linear combination of matrices of the form $\mathbf{x} \mathbf{x}^T$, $\mathbf{x} \in \mathbb{R}^n$.

If \mathbf{A} possesses any of these properties, it is said to be *positive semidefinite*. Decompositions of \mathbf{A} as a product $\mathbf{L}\mathbf{L}^T$ with a lower triangular square matrix \mathbf{L} is called a *Cholesky factorization* or *Cholesky decomposition* of \mathbf{A} .

The set of $n \times n$ real symmetric matrices is denoted by \mathbb{S}^n , the set of $n \times n$ positive semidefinite real symmetric matrices is denoted by \mathbb{S}_+^n . To avoid cumbersome notation, the superscript is dropped whenever the dimension is clear from the context. For the same reason, the vector order corresponding to \mathbb{S}_+^n (see Definition A.2) is simply denoted by \succcurlyeq , rather than $\succcurlyeq_{\mathbb{S}_+^n}$.

Using either definition from Proposition A.6 we see that \mathbb{S}_+^n is a convex cone. A few of its key properties are given in the following proposition; these are also direct consequences of Proposition A.6.

Proposition A.7. *The set of $n \times n$ positive semidefinite matrices is a closed, convex, and pointed cone. When embedded in \mathbb{S}^n , it is also a proper cone. Its boundary consists of the rank one matrices in the cone.*

The dual cone of \mathbb{S}_+^n can also be derived using Proposition A.6.

Proposition A.8. *When embedded in \mathbb{S}^n , $(\mathbb{S}_+^n)^* = \mathbb{S}_+^n$. When embedded in $\mathbb{R}^{n \times n}$, $(\mathbb{S}_+^n)^* = \mathbb{S}_+^n + \{\mathbf{A} \in \mathbb{R}^{n \times n} \mid \mathbf{A} = -\mathbf{A}^T\}$. (In the second equation the symbol $+$ denotes the Minkowski sum of sets.)*

A.3 Semidefinite and second-order cone programming

A *semidefinite optimization problem* or *semidefinite program* (SDP for short) is a problem of the form (A.1) with $\mathcal{K} = \mathbb{S}_+^n$. A *second-order conic program* (or SOCP for short) is a problem of the form (A.1) in which \mathcal{K} is the direct product of *second-order cones*. The $n + 1$ -dimensional second-order cone (also known as Lorentz cone, or simply ice cream cone) is the cone

$$\mathcal{Q}_{n+1} \stackrel{\text{def}}{=} \{(x_0, \mathbf{x}) \in \mathbb{R} \times \mathbb{R}^n \mid x_0 \geq \|\mathbf{x}\|\}.$$

Since both \mathcal{Q}_{n+1} and \mathbb{S}_+^n are self-dual, the duals of SOCPs and SDPs are also SOCPs and SDPs, respectively.

Semidefinite representable problems. Using a theorem of Haynsworth [Hay68] on the Schur complement of positive semidefinite matrices it is not hard to show that $(x_0, \mathbf{x}) \in \mathcal{Q}_{n+1}$ if and only if the corresponding *arrow-shaped matrix* $\begin{pmatrix} x_0 & \mathbf{x}^T \\ \mathbf{x} & \cdot \\ & & x_0 \end{pmatrix}$ is positive semidefinite. Thus, SDP generalizes SOCP, which in turn generalizes convex quadratic (and linear) programming. Nevertheless, special purpose algorithms developed for SOCP outperform the SDP algorithms on SOCPs in practice.

An optimization problem is *second-order cone representable* (or *semidefinite representable*) if, possibly by the introduction of new variables, it can be written equivalently as an SOCP (or SDP, respectively). For conic optimization problems this translates to the representability of the underlying cone \mathcal{K} by linear matrix inequalities.

Definition A.9. We say that a set X is *semidefinite representable* (or *SD-representable*) if

$$X = \{x \in \mathbb{R}^n \mid \exists u \in \mathbb{R}^k \text{ satisfying } C(x) + D(u) \succcurlyeq 0\}$$

for some affine mappings C, D with range in \mathbb{S}^m .

In particular, affine images and affine pre-images of positive semidefinite matrices are semidefinite representable. Furthermore, intersections and Minkowski sums of semidefinite representable sets are also semidefinite representable.

There is a vast literature on SOC- and SD-representable sets and functions, for an overview of the most important results the reader is referred to [BTN01, Chapters 3 and 4].

Algorithms and implementations. Both \mathcal{Q}_{n+1} and \mathbb{S}_+^n are convex cones with known self-concordant barrier functions for them. Hence, most of the standard interior-point methods [NN94, Ren01] are applicable to them. Furthermore, both cones are *symmetric*, that is, they are self-dual and homogeneous (their automorphism group acts transitively on their interior) [FK94]. This rich algebraic structure has been exploited in the development of algorithms for SOCPs and SDPs, starting with the interior point methods of [Fay97]. See also [AS00] for a survey on word-for-word translations of linear programming algorithms to SOCPs and SDPs, and [AG03] for efficient and numerically stable algorithms for SOCP. The theoretical complexity of optimization problems involving only linear and second-order cone constraints is the same as that of linear optimization [AG03].

Currently a number of implementations of various algorithms are available to solve SOCPs and SDPs. Many general purpose nonlinear solvers support SOCP constraints, including KNITRO [Zie], which the author used for most of the numerical experiments in this thesis.

SDPs are more difficult, and most optimization modeling languages and convex optimization solvers do not support semidefinite constraints. However, a number of Matlab [Mat] packages are available for solving SDPs, the ones most commonly used in the literature are Sedumi [Stu99] and SDPT3 [TTT99]. (Both of these are free, except for the fact that they require the commercial software Matlab.) The packages YALMIP [Löf04] and CVX [GB07] are designed to help translating SD-representable optimization problems to the standard form input these solvers require. CSDP

[Bor99] is an open source C library (now part of the COIN-OR project [LH03]) implementing a predictor-corrector interior point algorithm. SDPA [YFK03] is a multi-platform SDP solver designed specifically with sparse problems in mind. A common disadvantage of all of these solvers is that they can only handle problems that can be equivalently written in the standard form (A.1), but cannot handle general (not SD-representable) convex constraints.

Bibliography

- [AC10] Sylvain Arlot and Alain Celisse. A survey of cross-validation procedures for model selection. *Statistics Surveys*, 4:40–79, 2010. doi:10.1214/09-SS054. → p. 45.
- [AENR08] Farid Alizadeh, Jonathan Eckstein, Nilay Noyan, and Gábor Rudolf. Arrival rate approximation by nonnegative cubic splines. *Operations Research*, 56(1):140–156, Jan 2008. doi:10.1287/opre.1070.0443. → pp. 2, 25, 44, 53, and 68.
- [AG03] Farid Alizadeh and Donald Goldfarb. Second-order cone programming. *Mathematical Programming Series B*, 95(1):3–51, Jan 2003. doi:10.1007/s10107-002-0339-5. → pp. 33, 71, and 118.
- [Ali06] Farid Alizadeh. Semidefinite and second order cone programming and their application to shape constrained regression and density estimation. In *Proceedings of the INFORMS Annual Meeting*. INFORMS, Pittsburg, PA, 2006. → p. 9.
- [ANR05] Farid Alizadeh, Nilay Noyan, and Gábor Rudolf. Optimality constraints for the cone of positive polynomials. Tech. Rep. RRR 1–2005, Rutgers Center for Operations Research (RUTCOR), Rutgers University, Jan 2005. http://rutcor.rutgers.edu/pub/rrr/reports2005/1_2005.pdf. → p. 89.
- [Art27] Emil Artin. Über die Zerlegung definiter Funktionen in Quadrate. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 5(1):100–115, Dec 1927. doi:10.1007/BF02952513. → pp. 8 and 15.
- [AS00] Farid Alizadeh and Stefan H. Schmieta. Symmetric cones, potential reduction methods and word-by-word extensions. In Henry Wolkowicz, Romesh Saigal, and Lieven Vandenbergh (eds.), *Handbook of Semidefinite Programming, Theory, Algorithms and Applications*, pp. 195–233. Kluwer Academic Publishers, 2000. → pp. 89 and 118.
- [BA02] Kenneth P. Burnham and David R. Anderson. *Model Selection and Multimodel Inference*. 2nd ed. Springer, New York, NY, 2002. ISBN 0-387-95364-7. → pp. 29, 45, and 46.
- [BA04] Kenneth P. Burnham and David R. Anderson. Multimodel inference: Understanding AIC and BIC in model selection. *Sociological Methods & Research*, 33(2):261–304, Nov 2004. doi:10.1177/0049124104268644. → pp. 29 and 45.
- [BB05] Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. *Economic Theory*, 26(2):445–469, Aug 2005. doi:10.1007/s00199-004-0514-4. → p. 41.
- [BD94] Alain Berlinet and Luc Devroye. A comparison of kernel density estimates. *Publications de l'Institut de Statistique de l'Université de Paris*, 38(3):3–59, 1994. <http://cg.scs.carleton.ca/~luc/devs.html>. → p. 51.
- [BD07] Melanie Birke and Holger Dette. Estimating a convex function in nonparametric regression. *Scandinavian Journal of Statistics*, 34:384–404, Jun 2007. doi:10.1111/j.1467-9469.2006.00534.x. → p. 48.

- [Bec97] Bernhard Beckermann. The condition number of real Vandermonde, Krylov and positive definite Hankel matrices. *Numerische Mathematik*, 85(4):553–577, Jun 1997. doi:10.1007/PL00005392. → pp. 85 and 109.
- [Ben62] Jacques F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4:238–252, Dec 1962. doi:10.1007/BF01386316. → p. 63.
- [Ber99] Dimitri P. Bertsekas. *Nonlinear Programming*. 2nd ed. Athena Scientific, 1999. ISBN 978-1886529007. → p. 63.
- [BF84] V. Braibant and C. Fleury. Shape optimal design using B-splines. *Computer Methods in Applied Mechanics and Engineering*, 44(3):247–267, Aug 1984. doi:10.1016/0045-7825(84)90132-4. → p. 2.
- [BH02] Endre Boros and Peter L. Hammer. Pseudo-Boolean optimization. *Discrete Applied Mathematics*, 123(1–3):155–225, Nov 2002. doi:10.1016/S0166-218X(01)00341-9. → p. 15.
- [Boo01] Carl De Boor. *A practical guide to splines*, *Applied Mathematical Sciences*, vol. 27. revised ed. Springer, New York, NY, 2001. ISBN 0-387-95366-3. → pp. 24 and 37.
- [Bor99] Brian Borchers. CSDP, a C library for semidefinite programming. *Optimization Methods and Software*, 11(1):613–623, 1999. doi:10.1080/10556789908805765. → p. 119.
- [Bru55] H. D. Brunk. Maximum likelihood estimates of monotone parameters. *The Annals of Mathematical Statistics*, 26(4):607–616, Dec 1955. doi:10.1214/aoms/1177728420. → p. 1.
- [Bru58] H. D. Brunk. On the estimation of parameters restricted by inequalities. *Annals of Mathematical Statistics*, 29(2):437–454, Jun 1958. doi:10.1214/aoms/1177706621. → pp. 22, 23, and 28.
- [BS62] Louis Brickman and Leon Steinberg. On nonnegative polynomials. *The American Mathematical Monthly*, 69(3):218–221, Mar 1962. doi:10.2307/2311058. → pp. 11 and 12.
- [BTN01] Aharon Ben-Tal and Arkadi Nemirovski. *Lectures on Modern Convex Optimization*. SIAM, Philadelphia, PA, 2001. ISBN 0-89871-491-5. → p. 118.
- [Bur98] Christopher J. C. Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2(2):121–167, Jun 1998. doi:10.1023/A:1009715923555. → p. 82.
- [CLO04] David A. Cox, John Little, and Donal O’Shea. *Using Algebraic Geometry*. 2nd ed. Springer, New York, NY, 2004. ISBN 0-387-20706-6. → p. 14.
- [CVBM02] Olivier Chapelle, Vladimir Vapnik, Olivier Bousquet, and Sayan Mukherjee. Choosing multiple parameters for support vector machines. *Machine Learning*, 46(1-3):131–159, 2002. doi:10.1023/A:1012450327387. → p. 82.
- [d’A04] Alexandre d’Aspremont. *Shape Constrained Optimization, with Applications in Finance and Engineering*. Ph.D. thesis, Stanford University, Stanford, CA, 2004. → p. 23.
- [DG85] Luc Devroye and László Györfi. *Nonparametric Density Estimation – the L_1 view*. John Wiley & Sons, 1985. ISBN 0-471-81646-9. → pp. 1 and 51.
- [Dia62] P. H. Diananda. On non-negative forms in real variables some or all of which are non-negative. *Mathematical Proceedings of the Cambridge Philosophical Society*, 58:17–25, 1962. doi:10.1017/S0305004100036185. → p. 17.

- [Die80] P. Dierckx. Algorithm 42. an algorithm for for cubic spline fitting with convexity constraints. *Computing*, 24(4):349–371, Dec 1980. doi:10.1007/BF02237820. → p. 22.
- [DM61] M. L. Dudzinski and R. Mykutowycz. The eye lens as an indicator of age in the wild rabbit in Australia. *CSIRO Wildlife Research*, 6(2):156–159, 1961. doi:10.1071/CWR9610156. → pp. 48, 49, and 50.
- [DR03] Darinka Dentcheva and Andrzej Ruszczyński. Optimization with stochastic dominance constraints. *SIAM Journal on Optimization*, 14(2):548–566, 2003. doi:10.1137/S1052623402420528. → p. 3.
- [DR09] Lutz Dümbgen and Kaspar Rufibach. Maximum likelihood estimation of a log-concave density and its distribution function: Basic properties and uniform consistency. *Bernoulli*, 15(1):40–68, Feb 2009. doi:10.3150/08-BEJ141. → p. 41.
- [DS97] Holger Dette and William J. Studden. *The Theory of Canonical Moments with Applications in Statistics, Probability, and Analysis*. Wiley Interscience, New York, NY, Sep 1997. ISBN 0-471-10991-6. → pp. 10, 18, and 20.
- [DSTA96] Michel Delecroix, Michel Simioni, and Christine Thomas-Agnan. Functional estimation under shape constraints. *Journal of Nonparametric Statistics*, 6(1):69–89, 1996. doi:10.1214/aos/1009210689. → p. 24.
- [DTA00] Michel Delecroix and Christine Thomas-Agnan. Spline and kernel regression under shape restrictions. In Michael G. Schimek (ed.), *Smoothing and Regression*, chap. 5, pp. 109–134. Wiley-Interscience, 2000. → pp. 22 and 23.
- [DW60] George B. Dantzig and Philip Wolfe. Decomposition principle for linear programs. *Operations Research*, 8(1):101–111, Jan 1960. doi:10.1287/opre.8.1.101. → p. 63.
- [EL01] Paul P.B. Eggermont and Vincent N. LaRiccia. *Maximum Penalized Likelihood Estimation, Volume I: Density Estimation*. Springer-Verlag, New York, NY, 2001. ISBN 0-387-95268-3. → pp. 1, 23, 27, 39, 51, and 52.
- [Fay97] Leonid Faybusovich. Euclidean Jordan algebras and interior-point algorithms. *Positivity*, 1(4):331–357, Dec 1997. doi:10.1023/A:1009701824047. → pp. 89 and 118.
- [Fay02] Leonid Faybusovich. Self-concordant barriers for cones generated by Chebyshev systems. *SIAM Journal on Optimization*, 12(3):770–781, 2002. doi:10.1137/S1052623401386782. → p. 85.
- [FGK02] Robert Fourer, David M. Gay, and Brian W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. 2nd ed. Brooks/Cole–Thomson Learning, Pacific Grove, CA, 2002. ISBN 0-534-38809-4. → pp. 46 and 69.
- [FHT06] Tadayoshi Fushiki, Shingo Horiuchi, and Takashi Tsuchiya. A maximum likelihood approach to density estimation with semidefinite programming. *Neural Computation*, 18(11):2777–2812, Nov 2006. doi:10.1162/neco.2006.18.11.2777. → p. 25.
- [FK94] Jacques Faraut and Adam Korányi. *Analysis on Symmetric Cones*. Oxford Mathematical Monographs, Oxford University Press, Oxford, UK, 1994. ISBN 978-0-19-853477-8. → pp. 71, 89, and 118.
- [FMT08] Leonid Faybusovich, Thanasak Mouktonglang, and Takashi Tsuchiya. Numerical experiments with universal barrier functions for cones of Chebyshev systems. *Computational Optimization and Applications*, 41(2):205–223, Nov 2008. doi:10.1007/s10589-007-9101-5. → p. 85.

- [GB07] Michael Grant and Stephen Boyd. *CVX: Matlab software for disciplined convex programming (web page and software)*. Stanford University, Dec 2007. <http://cvxr.com/cvx/>.
→ pp. 32 and 118.
- [GG71] I. J. Good and R. A. Gaskins. Nonparametric roughness penalties for probability densities. *Biometrika*, 58(2):255–277, Aug 1971. doi:10.1093/biomet/58.2.255. → p. 23.
- [GH82] Stuart Geman and Chii-Ruey Hwang. Nonparametric maximum likelihood estimation by the method of sieves. *The Annals of Statistics*, 10(2):401–414, 1982. doi:10.1214/aos/1176345782. → p. 28.
- [GHND03] Y. Genin, Y. Hachez, Yu. Nesterov, and P. van Dooren. Optimization problems over positive pseudopolynomial matrices. *SIAM Journal on Matrix Analysis and Applications*, 25(1):57–79, 2003. doi:10.1137/S0895479803374840. → p. 80.
- [GL98] Miguel A. Goberna and Marco A. López. *Linear Semi-Infinite Optimization*. John Wiley & Sons, Chichester, UK, Mar 1998. ISBN 0-471-97040-9. → p. 3.
- [Gre81] Ulf Grenander. *Abstract Inference*. John Wiley & Sons, 1981. ISBN 0-471-08267-8. → p. 28.
- [GSY94] John A. Gregory, Muhammad Sarfraz, and P.K. Yuen. Interactive curve design using C^2 rational splines. *Computers & Graphics*, 18(2):153–159, Mar 1994. doi:10.1016/0097-8493(94)90089-2. → p. 2.
- [Hay68] Emilie V. Haynsworth. Determination of the inertia of a partitioned Hermitian matrix. *Linear Algebra and its Applications*, 1(1):73–81, Jan 1968. doi:10.1016/0024-3795(68)90050-5. → p. 117.
- [Hes69] Magnus R. Hestenes. Multiplier and gradient methods. *Journal of Optimization Theory and Applications*, 4:303–320, Nov 1969. doi:10.1007/BF00927673. → p. 64.
- [Hil88] David Hilbert. Ueber die Darstellung definiter Formen als Summe von Formenquadraten. *Mathematische Annalen*, 32:342–350, Sep 1888. doi:10.1007/BF01443605. → p. 15.
- [Hil00] David Hilbert. Mathematische Probleme. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen*, pp. 253–297, 1900. Also in *Archiv der Mathematik und Physik*, 3:44–63 and 213–237, 1901. English translation in *Bulletin of the AMS* 8:437–479 (1902). → pp. 8 and 15.
- [Hil54] Clifford Hildreth. Point estimates of ordinates of concave functions. *Journal of the American Statistical Association*, 49(267):598–619, Sep 1954. doi:10.2307/2281132, corrigenda doi:10.2307/2281222. → pp. 1 and 22.
- [HJ85] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985. ISBN 0-521-30586-1. → pp. 114 and 116.
- [HJ90] Georg Heinig and P. Jankowski. Parallel and superfast algorithms for Hankel systems of equations. *Numerische Mathematik*, 58(1):109–127, Dec 1990. doi:10.1007/BF01385613. → p. 108.
- [HL03] Didier Henrion and Jean-Bernard Lasserre. GloptiPoly: Global optimization over polynomials with Matlab and SeDuMi. *ACM Transactions on Mathematical Software*, 29(2):165–194, Jun 2003. doi:10.1145/779359.779363. → p. 8.
- [HLP34] Godfrey Harold Hardy, John Edensor Littlewood, and George Pólya. *Inequalities*. Cambridge University Press, 1934. → pp. 15 and 16.
- [HR69] Josef Hadar and William R. Russell. Rules for ordering uncertain prospects. *American Economic Review*, 59(1):25–34, Mar 1969. → p. 2.

- [HS98] Xuming He and Peide Shi. Monotone B-spline smoothing. *Journal of the American Statistical Association*, 93(442):643–650, Jun 1998. doi:10.2307/2670115. → p. 22.
- [HUS10] Jean-Baptiste Hiriart-Urruty and Alberto Seeger. A variational approach to copositive matrices. *SIAM Review*, 52(4):593–629, Dec 2010. doi:10.1137/090750391. → p. 16.
- [ILO] ILOG S.A. ILOG CPLEX documentation, version 11.0. <http://www.decf.berkeley.edu/help/apps/AMPL/cplex-doc/>. → pp. 33 and 46.
- [Kar63] Samuel Karlin. Representation theorems for positive functions. *Journal of Mathematics and Mechanics (Indiana University Mathematics Journal)*, 12:599–617, 1963. doi:10.1512/iumj.1963.12.12038. → p. 13.
- [KLP05] Etienne de Klerk, Monique Laurent, and Pablo Parrilo. On the equivalence of algebraic approaches to the minimization of forms on the simplex. In Didier Henrion and Andrea Garulli (eds.), *Positive Polynomials in Control, Lecture Notes in Control and Information Sciences*, vol. 312, pp. 121–132. Springer, Berlin/Heidelberg, 2005. ISBN 3-540-23948-0. doi:10.1007/10997703_7. → p. 17.
- [KM07] Masakazu Kojima and Masakazu Muramatsu. An extension of sums of squares relaxations to polynomial optimization problems over symmetric cones. *Mathematical Programming Series A*, 110(2):315–336, Jul 2007. doi:10.1007/s10107-006-0004-5. → pp. 71 and 82.
- [KMY03] Piyush Kumar, Joseph S. B. Mitchell, and E. Alper Yildirim. Computing core-sets and approximate smallest enclosing hyperspheres in high dimensions. In Richard E. Ladner (ed.), *Proceedings of the Fifth Workshop on Algorithm Engineering and Experiments*, pp. 45–55. SIAM, Baltimore, MD, 2003. ISBN 0898715423. → p. 82.
- [KN63] John L. Kelley and Isaac Namioka. *Linear Topological Spaces*. No. 36 in Graduate Texts in Mathematics, Springer, New York, NY, 1963. ISBN 0-387-90169-8. → p. 114.
- [KN77] Mark Grigor’evich Kreĭn and Adol’f Abramovich Nudel’man. *The Markov Moment Problem and Extremal Problems*. AMS, Providence, RI, 1977. ISBN 0-8218-4500-4. → p. 18.
- [Koc04] Jerzy Kociński. *Cracovian Algebra*. Nova Science, 2004. ISBN 1-59454-105-1. → p. 76.
- [Koe99] Max Koecher. *The Minnesota Notes on Jordan Algebras and Their Applications, Lecture Notes in Mathematics*, vol. 1710. Springer-Verlag, 1999. ISBN 3-540-66360-0. → p. 89.
- [Koj03] Masakazu Kojima. Sums of squares relaxations of polynomial semidefinite programs. Tech. Rep. B-397, Dept. of Math and Comp. Sci, Tokyo Institute of Technology, Nov 2003. <http://www.is.titech.ac.jp/research/research-report/B/B-397.ps.gz>. → p. 71.
- [KRR99] Krzysztof C. Kiwiel, Charles H. Rosa, and Andrzej Ruszczyński. Decomposition in large system optimization using the method of multipliers. *SIAM Journal on Optimization*, 9:668–689, 1999. doi:10.1137/S1052623495288064. → p. 64.
- [KS53] Samuel Karlin and Lloyd S. Shapley. *Geometry of Moment Spaces, Memoirs of the AMS*, vol. 12. AMS, Providence, RI, 1953. ISBN 9780608201542. → p. 13.
- [KS66] Samuel Karlin and William J. Studden. *Tchebycheff Systems, with Applications in Analysis and Statistics, Pure and Applied Mathematics*, vol. XV. Wiley Interscience, New York, NY, 1966. → pp. 10, 12, 20, 32, 100, 101, and 102.
- [KW71] George Kimeldorf and Grace Wahba. Some results on Tchebycheffian spline functions. *Journal of Mathematical Analysis and Applications*, 33(1):82–95, Jan 1971. doi:10.1016/0022-247X(71)90184-3. → p. 27.

- [Lan87] Henry J. Landau. *Moments in Mathematics, Proceedings of Symposia in Applied Mathematics*, vol. 37. AMS, 1987. ISBN 0-8218-0114-7. → p. 18.
- [Las70] Leon S. Lasdon. *Optimization Theory for Large Systems*. Macmillan, 1970. ISBN 978-0023678004. Also available from Dover (2002). → p. 63.
- [Las01a] Jean B. Lasserre. An explicit exact SDP relaxation for nonlinear 0-1 programs. In Karen Aardal and A. M. H. Gerards (eds.), *Integer Programming and Combinatorial Optimization*, pp. 293–303. No. 2081 in Lecture Notes in Computer Science, Springer, 2001. doi:10.1007/3-540-45535-3.23. → p. 84.
- [Las01b] Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001. doi:10.1137/S1052623400366802. → p. 17.
- [Lau03] Monique Laurent. A comparison of the Sherali–Adams, Lovász–Schrijver, and Lasserre relaxations for 0-1 programming. *Mathematics of Operations Research*, 28(3):470–496, Aug 2003. doi:10.1287/moor.28.3.470.16391. → p. 84.
- [Lau08] Monique Laurent. Sums of squares, moment matrices, and optimization over polynomials. In Mihai Putinar and Seth Sullivant (eds.), *Emerging Applications of Algebraic Geometry, IMA Volumes in Mathematics and its Applications*, vol. 149, pp. 157–270. Springer, 2008. ISBN 978-0-387-09685-8. → p. 7.
- [LH03] Robin Lougee-Heimer. The Common Optimization INterface for Operations Research. *IBM Journal of Research and Development*, 47(1):57–66, Jan 2003. doi:10.1147/rd.471.0057. → p. 119.
- [Li97] Tien-Yien Li. Numerical solution of multivariate polynomial systems by homotopy continuation methods. *Acta Numerica*, 6:399–436, 1997. doi:10.1017/S0962492900002749. → p. 14.
- [Löf04] Johan Löfberg. YALMIP: A toolbox for modeling and optimization in MATLAB. In *Proceedings of the CACSD Conference*. Taipei, Taiwan, 2004. URL <http://users.isy.liu.se/johanl/yalmip>, doi:10.1109/CACSD.2004.1393890. → pp. 8 and 118.
- [Lor86] G. G. Lorentz. *Bernstein Polynomials*. 2nd ed. AMS Chelsea Publishing, New York, NY, 1986. ISBN 978-0-8284-0323-8. → p. 30.
- [LS78] Peter A.W. Lewis and Gerald S. Shedler. Simulation methods for Poisson processes in nonstationary systems. In Harold Joseph Highland (ed.), *Proceedings of the 10th conference on Winter simulation*, pp. 155–163. IEEE Press, Piscataway, NJ, Dec 1978. → p. 53.
- [Luk18] Franz Lukács. Verschärfung der ersten Mittelwertsatzes der Integralrechnung für rationale Polynome. *Mathematische Zeitschrift*, 2:229–305, 1918. doi:10.1007/BF01199412. → p. 11.
- [Mar08] Murray Marshall. *Positive Polynomials and Sums of Squares*. AMS, 2008. ISBN 978-0-8218-4402-1. → pp. 7 and 15.
- [Mat] MathWorks. MATLAB documentation. <http://www.mathworks.com/help/techdoc/>. → pp. 8 and 118.
- [Mey96] Mary Catherine Meyer. *Shape Restricted Inference with Applications to Nonparametric Regression, Smooth Nonparametric Function Estimation, and Density Estimation*. Ph.D. thesis, University of Michigan, Ann Arbor, MI, 1996. → pp. 22 and 23.
- [MH03] John C. Mason and David C. Handscomb. *Chebyshev Polynomials*. CRC Press, Boca Raton, FL, 2003. ISBN 0-8493-0355-9. → pp. 103, 104, and 110.

- [MLB05] Alessandro Magnani, Sanjay Lall, and Stephen Boyd. Tractable fitting with convex polynomials via sum-of-squares. In *44th IEEE Conference on Decision and Control and 2005 European Control Conference. CDC-ECC '05.*, pp. 1672–1677. Seville, Spain, Dec 2005. doi:10.1109/CDC.2005.1582399. → p. 70.
- [MNP04] Gergely Mádi-Nagy and András Prékopa. On multivariate discrete moment problems and their applications to bounding expectations and probabilities. *Mathematics of Operations Research*, 29(2):229–258, May 2004. doi:10.1287/moor.1030.0064. → p. 42.
- [Mot67] Theodore Samuel Motzkin. The arithmetic-geometric inequality. In *Inequalities (Proc. Sympos. Wright-Patterson Air Force Base, Ohio, 1965)*, pp. 205–224. Academic Press, New York, NY, 1967. → p. 14.
- [MP89] Charles A. Micchelli and Allan Pinkus. Some remarks on nonnegative polynomials on polyhedra. In T.W. Anderson, Krishna B. Athreya, and Donald L. Iglehart (eds.), *Probability, Statistics, and Mathematics: Papers in honor of Samuel Karlin*, pp. 163–186. Academic Press, San Diego, CA, 1989. ISBN 0-12-058470-0. → p. 17.
- [MTA99] E. Mammen and C. Thomas-Agnan. Smoothing splines and shape restrictions. *Scandinavian Journal of Statistics*, 26(2):239–252, Jun 1999. doi:10.1111/1467-9469.00147. → pp. 24 and 25.
- [MTT75] G. F. de Montricher, R. A. Tapia, and J. R. Thompson. Nonparametric maximum likelihood estimation of probability densities by penalty function methods. *The Annals of Statistics*, 3(6):1329–1348, Nov 1975. doi:10.1214/aos/1176343288. → pp. 22 and 23.
- [Nes97] Yurii Nesterov. Structure of non-negative polynomials and optimization problems, 1997. Unpublished manuscript. → pp. 102 and 103.
- [Nes00] Yurii Nesterov. Squared functional systems and optimization problems. In Hans Frenk, Kees Roos, Tamás Terlaky, and Shuzhong Zhang (eds.), *High Performance Optimization, Applied Optimization*, vol. 33, pp. 405–440. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2000. ISBN 0-7923-6013-3. → pp. 9, 12, 58, 70, 71, 75, and 78.
- [NN94] Yurii Nesterov and Arkadii Nemirovski. *Interior-Point Polynomial Algorithms in Convex Programming, Studies in Applied Mathematics*, vol. 13. SIAM, Philadelphia, PA, 1994. ISBN 0-89871-319-6. → pp. 85 and 118.
- [NPT84] A. S. Nemirovskii, B. T. Polyak, and A. B. Tsybakov. Signal processing by the nonparametric maximum-likelihood method. *Problemy peredachi informatsii*, 20(3):29–46, 1984. English translation in *Problems of Information Transmission* 20:177-192 (1984). → pp. 22, 27, and 28.
- [NPT85] A. S. Nemirovskii, B. T. Polyak, and A. B. Tsybakov. Convergence rate of nonparametric estimates of maximum-likelihood type. *Problemy peredachi informatsii*, 21(4):17–33, 1985. English translation in *Problems of Information Transmission* 21:258-272 (1985). → p. 22.
- [NS07] Jiawang Nie and Markus Schweighofer. On the complexity of Putinar’s Positivstellensatz. *Journal of Complexity*, 23(1):135–150, Feb 2007. doi:10.1016/j.jco.2006.07.002. → p. 17.
- [OzD] OzDASL. Australasian Data and Story Library. <http://www.statsci.org/data/index.html>. → p. 48.
- [PA] Dávid Papp and Farid Alizadeh. Linear and second order cone programming approaches to statistical estimation problems. *Under review in Statistics and Computing*. → p. 6.

- [PA08] Dávid Papp and Farid Alizadeh. Linear and second order cone programming approaches to statistical estimation problems. Tech. Rep. RRR 13–2008, Rutgers Center for Operations Research (RUTCOR), Rutgers University, Sep 2008. http://rutcor.rutgers.edu/pub/rrr/reports2008/13_2008.pdf. → pp. 2 and 6.
- [Pan01] Victor Y. Pan. *Structured matrices and polynomials: unified superfast algorithms*. Birkhäuser, Boston, MA, 2001. ISBN 0-8176-4240-4. → pp. 108 and 110.
- [Par00] Pablo A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. Ph.D. thesis, California Institute of Technology, Pasadena, CA, May 2000. → p. 16.
- [Par03] Pablo A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96(2):293–320, 2003. doi:10.1007/s10107-003-0387-5. → p. 17.
- [PCA] Dávid Papp, Ricardo A. Collado, and Farid Alizadeh. Extension of the semidefinite characterization of sum of squares functional systems to algebraic structures. *Under review in Linear Algebra and Its Applications*. → p. 6.
- [PPSP04] Stephen Prajna, Antonis Papachristodoulou, Peter Seiler, and Pablo A. Parrilo. *SOSTOOLS: Sum of squares optimization toolbox for MATLAB*, 2004. <http://www.cds.caltech.edu/sostools>. → p. 8.
- [Pré71] András Prékopa. Logarithmic concave measures with application to stochastic programming. *Acta Scientiarum Mathematicarum*, 32(3–4):301–316, 1971. <http://rutcor.rutgers.edu/~prekopa/SCIENT1.pdf>. → p. 41.
- [Pré90] András Prékopa. The discrete moment problem and linear programming. *Discrete Applied Mathematics*, 27(3):235–254, Jun 1990. doi:10.1016/0166-218X(90)90068-N. → p. 42.
- [PS76] George Pólya and Gabor Szegő. *Problems and Theorems in Analysis*, vol. 2. Springer, New York, NY, 1976. ISBN 0-387-06972-0. → p. 10.
- [PS03] Pablo A. Parrilo and Bernd Sturmfels. Minimizing polynomial functions. In Saugata Basu and Laureano Gonzales-Vega (eds.), *Algorithmic and Quantitative Real Algebraic Geometry, DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, vol. 60, pp. 83–99. AMS, 2003. ISBN 0-8218-2863-0. → p. 14.
- [Put93] Mihai Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal*, 42:969–984, 1993. → p. 16.
- [PWM07] Jayanta Kumar Pal, Michael Woodroffe, and Mary Meyer. Estimating a Poly frequency function₂. In Regina Liu, William Strawderman, and Cun-Hui Zhang (eds.), *Complex datasets and inverse problems, IMS Lecture Notes Monogr. Ser.*, vol. 54, pp. 239–249. Institute of Mathematical Statistics, Beachwood, OH, 2007. doi:10.1214/074921707000000184. → p. 41.
- [Ram88] J. O. Ramsay. Monotone regression splines in action. *Statistical Science*, 3(4):425–441, 1988. doi:10.1214/ss/1177012761. → pp. 24, 25, and 30.
- [Ren92] James Renegar. On the computational complexity and geometry of the first-order theory of the reals. Parts I–III. *Journal of Symbolic Computation*, 13(3):255–352, Mar 1992. doi:10.1016/S0747-7171(10)80003-3. → pp. 14 and 57.
- [Ren01] James Renegar. *A Mathematical View of Interior-Point Methods in Convex Optimization*. No. 03 in MPS-SIAM Series on Optimization, MPS-SIAM, Philadelphia, PA, 2001. ISBN 0-89871-502-4. → pp. 86, 114, and 118.

- [RNPA] Gábor Rudolf, Nilay Noyan, Dávid Papp, and Farid Alizadeh. Bilinearity rank of the cone of positive polynomials and related cones. *Mathematical Programming*. Accepted. → pp. 6, 88, 90, and 91.
- [Roc70] Ralph Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1970. ISBN 0-691-08069-0. → p. 72.
- [Roc76] Ralph Tyrrell Rockafellar. Augmented Lagrangians and applications of the proximal point algorithm in convex programming. *Mathematics of Operations Research*, 1:97–116, May 1976. doi:10.1287/moor.1.2.97. → pp. 64 and 66.
- [RR98] Rembert Reemtsen and Jan J. Rückmann (eds.). *Semi-Infinite Programming*. Springer, Apr 1998. ISBN 978-0792350545. → p. 3.
- [Rud09] Gábor Rudolf. *New Results in Optimization with Functional Constraints*. Ph.D. thesis, Rutgers, The State University of New Jersey, New Brunswick, NJ, Oct 2009. <http://hdl.rutgers.edu/1782.2/rucore10001600001.ETD.000051899>. → pp. 2, 86, and 89.
- [Ruf07] Kaspar Rufibach. Computing maximum likelihood estimators of a log-concave density function. *Journal of Statistical Computation and Simulation*, 77(7–8):561–574, Jan 2007. doi:10.1080/10629360600569097. → p. 41.
- [Rus95] Andrzej Ruszczyński. On convergence of an augmented Lagrangian decomposition method for sparse convex optimization. *Mathematics of Operations Research*, 20:634–656, Aug 1995. doi:10.1287/moor.20.3.634. → pp. 63, 64, 65, and 66.
- [Rus05] Andrzej Ruszczyński. *Nonlinear Optimization*. Princeton University Press, Princeton, NJ, 2005. ISBN 978-0-691-11915-1. → p. 64.
- [RWD88] Tim Robertson, F.T. Wright, and R.L. Dykstra. *Order restricted statistical inference*. Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons, Chichester, UK, 1988. ISBN 0471917877. → pp. 1, 22, and 23.
- [Sch73] E. J. Schlossmacher. An iterative technique for absolute deviations curve fitting. *Journal of the American Statistical Association*, 68(344):857–859, Dec 1973. doi:10.2307/2284512. → p. 49.
- [Sch81] Larry L. Schumaker. *Spline Functions: Basic Theory*. Wiley-Interscience, 1981. ISBN 0-471-76475-2. → pp. 4, 30, 35, and 60.
- [Sch91] Konrad Schmüdgen. The K -moment problem for compact semi-algebraic sets. *Mathematische Annalen*, 289:203–206, 1991. doi:10.1007/BF01446568. → p. 16.
- [Sch00] Michael G. Schimek (ed.). *Smoothing and Regression*. John Wiley & Sons, 2000. ISBN 978-0-471-17946-7. → p. 23.
- [Sch04] Markus Schweighofer. On the complexity of Schmüdgen’s Positivstellensatz. *Journal of Complexity*, 20(4):529–543, Aug 2004. doi:10.1016/j.jco.2004.01.005. → p. 17.
- [Sco76] David W. Scott. *Nonparametric probability density estimation by optimization theoretic techniques*. Ph.D. thesis, Rice University, Houston, TX, 1976. → p. 22.
- [SSP09] Ersoy Subasi, Mine Subasi, and András Prékopa. Discrete moment problems with distributions known to be unimodal. *Mathematical Inequalities & Applications*, 12(3):587–610, Jul 2009. → p. 42.
- [ST43] James Alexander Shohat and Jacob David Tamarkin. *The problem of moments, Mathematical Surveys and Monographs*, vol. 1. AMS, 1943. ISBN 978-0-8218-1501-4. → pp. 19 and 20.

- [Str88] Dirk J. Struik. *Lectures on Classical Differential Geometry*. Dover, New York, NY, 1988. ISBN 0-486-65609-8. → p. 71.
- [STT80] David W. Scott, Richard A. Tapia, and James R. Thompson. Nonparametric probability density estimation by discrete maximum penalized-likelihood criteria. *The Annals of Statistics*, 8(4):820–832, 1980. doi:10.1214/aos/1176345074. → p. 22.
- [Stu99] Jos F. Sturm. Using SeDuMi 1.02, a Matlab toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11–12(1–4):625–653, 1999. doi:10.1080/10556789908805766. → pp. 8, 33, and 118.
- [Tar51] Alfred Tarski. A decision method for elementary algebra and geometry. Tech. Rep. R-109, RAND Corporation, May 1951. <http://www.rand.org/pubs/reports/2008/R109.pdf>. → p. 14.
- [TT90] James R. Thompson and Richard A. Tapia. *Nonparametric Function Estimation, Modeling, and Simulation*. SIAM, Philadelphia, PA, 1990. ISBN 0-89871-261-0. → pp. 22 and 25.
- [TTT99] Kim Chuan Toh, Michael J. Todd, and R. H. Tütüncü. SDPT3 — a Matlab software package for semidefinite programming, version 1.3. *Optimization Methods and Software*, 11–12(1–4):545–581, 1999. doi:10.1080/10556789908805762. → pp. 8 and 118.
- [TTT03] Reha H. Tütüncü, Kim-Chuan Toh, and Michel J. Todd. Solving semidefinite-quadratic-linear programs using SDPT3. *Mathematical Programming Series B*, 95(2):189–217, Feb 2003. doi:10.1007/s10107-002-0347-5. → p. 33.
- [Tur05] Berwin A. Turlach. Shape constrained smoothing using smoothing splines. *Computational Statistics*, 20(1):81–103, Mar 2005. doi:10.1007/BF02736124. → pp. 22 and 24.
- [Tyr94] Evgenij E. Tyrtyshnikov. How bad are Hankel matrices? *Numerische Mathematik*, 67(2):261–269, Mar 1994. doi:10.1007/s002110050027. → pp. 85 and 109.
- [Van06] Robert J. Vanderbei. LOQO user’s manual. Tech. rep., Operations Research and Financial Engineering, Princeton University, Sep 2006. <http://www.princeton.edu/~rvdb/tex/loqo/loqo405.pdf>. → p. 33.
- [VW87] Miguel Villalobos and Grace Wahba. Inequality-constrained multivariate smoothing splines with application to the estimation of posterior probabilities. *Journal of the American Statistical Association*, 82(397):239–248, 1987. doi:10.2307/2289160. → pp. 24, 27, 57, 61, and 62.
- [Wah90] Grace Wahba. *Spline Models for Observational Data*. No. 59 in CBMS-NSF Regional Conference series in applied mathematics, SIAM, Philadelphia, PA, 1990. ISBN 0-89871-244-0. → pp. 4, 27, 45, and 63.
- [Wan78] Jude Wanniski. Taxes, revenues, and the “Laffer curve”. *The Public Interest*, pp. 3–16, 1978. → p. 1.
- [WL08] Xiao Wang and Feng Li. Isotonic smoothing spline regression. *Journal of Computational and Graphical Statistics*, 17(1):21–37, Mar 2008. doi:10.1198/106186008X285627. → p. 25.
- [WNM76] N. Watanabe, Y. Nishimura, and M. Matsubara. Decomposition in large system optimization using the method of multipliers. *Journal of Optimization Theory and Applications*, 25:181–193, Jun 1976. doi:10.1007/BF00933211. → p. 64.

- [WSV00] Henry Wolkowicz, Romesh Saigal, and Lieven Vandenbergh (eds.). *Handbook of Semidefinite Programming: Theory, Algorithms, and Applications*, *International series in operations research & management science*, vol. 27. Kluwer, Norwell, MA, 2000. ISBN 0-7923-7771-0. → p. 114.
- [YFK03] Makoto Yamashita, Katsuki Fujisawa, and Masakazu Kojima. Implementation and evaluation of SDPA 6.0 (Semidefinite Programming Algorithm 6.0). *Optimization Methods and Software*, 18(4):491–505, Aug 2003. doi:10.1080/1055678031000118482. → p. 119.
- [You61] D. C. Youla. On the factorization of rational matrices. *IRE Transactions on Information Theory*, pp. 172–189, Jul 1961. doi:10.1109/TIT.1961.1057636, corrigenda doi:10.1109/TIT.1961.1057666. → pp. 79 and 80.
- [Zie] Ziena. KNITRO documentation. <http://www.ziena.com/documentation.htm>. → pp. 33, 46, 69, and 118.

VITA

Dávid Papp

- 2005-2011 Rutgers University, New Brunswick, NJ, USA
 Ph.D., Operations Research, May 2011
 M.S., Operations Research, May 2007
- 2000-2005 Budapest University of Technology and Economics, Budapest, Hungary
 M.S., Computer Science, June 2005

Publications

Scenario decomposition of risk-averse multistage stochastic programming problems. Accepted in *Annals of Operations Research* in March 2011. (With R.A. Collado, A. Ruszczyński)

Acyclic, or totally tight, two-person game forms; characterization and main properties. *Discrete Mathematics* 310(6-7), April 2010, pp.1135-1151. (With E. Boros, V. Gurvich, K. Makino)

Bilinearity Rank of the Cone of Positive Polynomials and Related Cones. Accepted in *Mathematical Programming* in January 2010. (With N. Noyan, G. Rudolf, F. Alizadeh)

Effective solution of linear Diophantine equation systems with an application in chemistry. *Journal of Mathematical Chemistry* 39(1), January 2006, pp.15-31. (With B. Vizvari)

Analysis of Petri net-based Models and their Applications in Reaction Kinetics. (In Hungarian, with English summary.) M.Sc. thesis, May 2005

Visualization of Nil-geometry. *Proceedings of the Dresden Geometry Symposium*, November 2003. (With E. Molnar)

Hardware-software partitioning in embedded system design. *Proceedings of the 2003 IEEE International Symposium on Intelligent Signal Processing*, Budapest, 2003. (With P. Arato, S. Juhasz, Z. Mann, A. Orban)