CentER

Asymptotic Analysis of Semidefinite Bounds for Polynomial Optimization and Independent Sets in Geometric Hypergraphs

LUCAS SLOT

Tilburg University

Asymptotic Analysis of Semidefinite Bounds for Polynomial Optimization and Independent Sets in Geometric Hypergraphs

Proefschrift ter verkrijging van de graad van doctor aan Tilburg University op gezag van de rector magnificus, prof. dr. W.B.H.J. van de Donk, in het openbaar te verdedigen ten overstaan van een door het college voor promoties aangewezen commissie in de Aula van de Universiteit op

vrijdag 30 september 2022 om 10:00 uur

door

Lucas Frederik Herman Slot,

geboren te Amsterdam

Promotores:	prof. dr. M. Laurent (Tilburg University)
	prof. dr. E. de Klerk (Tilburg University)
Promotiecommissie:	prof. dr. ir. E.R. van Dam (Tilburg University)
	dr. H. Fawzi (University of Cambridge)
	prof. dr. D.C. Gijswijt (TU Delft)
	dr. J.B. Lasserre (LAAS-CNRS, IMT and ANITI)
	prof. dr. M. Schweighofer (University of Konstanz)
	prof. dr. ir. R. Sotirov (Tilburg University)
	dr. J.C. Vera Lizcano (Tilburg University)

This research was carried out at the Centrum Wiskunde & Informatica (CWI) in Amsterdam.
This research has received funding from the European Union's Framework Programme for Research and Innovation Horizon 2020 under the Marie Skłodowska-Curie Actions Grant Agreement No. 764759 (MINOA).

Acknowledgments

First, I would like to thank my advisors, Monique and Etienne. For introducing me to the world of academic research, for showing me the ropes, and for guiding me towards independency. Thanks Monique, for your extraordinary commitment and dedication, witnessed by many late night emails and long (usually productive) meetings.

I thank the members of my PhD committee, Edwin van Dam, Hamza Fawzi, Dion Gijswijt, Jean-Bernard Lasserre, Markus Schweighofer, Renata Sotirov, and Juan Vera Lizcano for their valuable comments on the thesis and many interesting questions.

Next, I want to thank my co-authors Davi, Fernando and Frank, for our discussions over Zoom, where we triangulated the contents of the final chapter of this thesis. A special thanks to Frank and Davi, for their hospitality during my stay in Cologne, and to Fernando, for teaching me the importance of non-breaking spaces in LaTeX.

I thank the many great people at CWI and in Tilburg for their company, both at work and away on conferences (the few that were not cancelled). Let me mention in particular Sander, for the games of foosball and table tennis. Sven, for the timely updates on the games of Ajax. Andries, for his endless reservoir of tall stories and interesting, if a little disturbing, facts. Luis Felipe, for trying to teach me the salsa at a bar in Toulouse. Daniel, for introducing me to the marvelous worlds of Brandon Sanderson. Felix, for his appreciation of good food and beer.

A special thanks goes to David, for our extended coffee breaks, during which we discussed at length the few things in life more important than mathematics, such as optimal drafting strategies in *Magic: The Gathering*.

I thank my paranymph Pepijn, with whom I have shared most of my mathematical journey. From the fellowship of the *dubbele bachelor* at Science Park, through the two towers of *Tannenbusch* in Bonn and finally our return to *De Pijp* in Amsterdam.

I thank my parents, Jeroen and Desirée, for their endless support. My sister and paranymph Emma, who meticulously planned my defense ceremony (and the accompanying party). Last, but not least, Fleur, for her love and patience.

Contents

Introduction

If we wait for the moment when everything, absolutely everything is ready, we shall never begin.

Ivan Turgenev, Fathers and Sons

A mathematical optimization problem asks us to maximize an objective (or minimize a cost) under a given set of constraints. Solving such problems has many applications; both in the real world, and in other fields of mathematics. Unfortunately, some of the most interesting optimization problems are very difficult to solve algorithmically. One way to work around this issue is to consider so-called *relaxations*. That is, to consider variants of the problem which are (much) easier to solve, but whose solutions still provide good *approximations* for the original problem. In this thesis, we look at relaxations to hard problems based on *semidefinite programming*. We can distinguish two settings.

First, we consider so-called *sum-of-squares hierarchies*. These hierarchies allow one to define increasingly accurate – but also more computationally expensive – relaxations for *polynomial* optimization problems. They yield very good approximations in practice, and as a result they have been widely applied and studied in the literature. The central question we wish to answer in this setting is whether we can back their good performance up with *theoretical guarantees*. We address this question in Part 1 and Part 2 of the thesis, which each focus on a particular type of hierarchy.

Second, we look at the problem of finding *independent sets* of maximum size in a graph G, which is a classical example of an NP-hard problem. The celebrated *Lovász theta number* ϑ bounds the *independence number* of G from above. It has been succesfully applied to the setting of *geometric graphs*, yielding new results in discrete geometry and extremal combinatorics. In Part 3 of the thesis, we will generalize ϑ to (geometric) *hyper*graphs. A careful analysis of the resulting approximations allows us to improve an existing result in *Euclidean Ramsey theory*.

As we shall see, the three parts that make up this thesis are rather connected. First off, both the sum-of-squares hierarchies of Parts 1 and 2 and the Lovász theta number of Part 3 are examples of semidefinite programs. In fact, ϑ may be viewed in some sense as the 'first level' of a particular sum-of-squares hierarchy. Beyond that, the methods we use to prove new results in each of these settings share many similarities. In short, they all rely on classical *orthogonal polynomials*, and the relation between these polynomials and *polynomial (or continuous) kernels*. By exploiting symmetry, these relations allow us to move from a difficult, *multivariate* setting to a simpler, *univariate* setting, where an *asymptotic* analysis is then possible. We explain these connections in more detail in Chapter 1.

Polynomial optimization and sums of squares

Let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree d in n variables. We say that f is a *sum of squares* if there exist polynomials $p_1, p_2, \ldots, p_\ell \in \mathbb{R}[\mathbf{x}]$ such that:

$$f(\mathbf{x}) = p_1(\mathbf{x})^2 + p_2(\mathbf{x})^2 + \ldots + p_\ell(\mathbf{x})^2.$$

Clearly, if f is a sum of squares, then it is globally nonnegative; meaning that $f(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$. A natural question – which goes back to work of Hilbert in the late 19th century – is whether the converse is also true. This turns out not to be the case: Hilbert shows that all nonnegative polynomials in n variables of degree d are sums of squares if and only if $d = 2$, $n = 1$, or $n = 2$ and $d = 4$. The first explicit example of a nonnegative polynomial which is not a sum of squares was given much later in 1967 by Motzkin: $f(\mathbf{x}) = \mathbf{x}_1^4\mathbf{x}_2^2 + \mathbf{x}_1^2\mathbf{x}_2^4 - 3\mathbf{x}_1^2\mathbf{x}_2^2 + 1$. In 1927, Artin showed that *any* nonnegative polynomial may be expressed as a sum of squares of rational functions, thereby solving Hilbert's 17th problem posed in 1900. Later results in real algebraic geometry show existence of *structured* sum-of-squares decompositions for *positive* polynomials on *semialgebraic sets* (cf. [**PD01**]). In recent decades, these *Positivstellensätze* have found a new application in the field of mathematical programming; more specifically in *polynomial optimization*.

A polynomial optimization problem asks to minimize a given polynomial $f \in \mathbb{R}[\mathbf{x}]$ over a semialgebraic set $\mathbf{X} \subseteq \mathbb{R}^n$, which is itself defined by polynomials $g_1, g_2, \ldots, g_m \in \mathbb{R}[\mathbf{x}]$:

$$f_{\min} := \min_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x}), \qquad \text{where } \mathbf{X} = \{\mathbf{x} \in \mathbb{R}^n : g_1(\mathbf{x}) \geq 0, \ldots, g_m(\mathbf{x}) \geq 0\}.$$

Polynomial optimization problems are very general. They naturally capture classical NP-hard combinatorial problems including MaxCut and StableSet, even when the feasible region $\mathbf{X}$ is a relatively simple set, such as the unit sphere, the unit ball, the binary hypercube or the standard simplex. Further applications are found in finance, energy optimization, machine learning, optimal control and quantum information theory.

Sum-of-squares hierarchies based on Positivstellensätze

In light of their broad applicability, polynomial optimization problems are unsurprisingly difficult to solve numerically. Often, the best one can do is to *approximate* the value of $f_{\min}$. Perhaps the most well-known and successful methods for computing such approximations are so-called *sum-of-squares hierarchies*, due to Lasserre [**Las01**] and Parillo [**Par00**] in the early 2000s. The key idea underlying these hierarchies is that nonnegativity of the polynomial f on the set $\mathbf{X}$ may be verified by finding sum-of-squares polynomials $\sigma_0, \sigma_1, \ldots \sigma_m \in \Sigma[\mathbf{x}]$ so that:

$$f(\mathbf{x}) = \sigma_0(\mathbf{x}) + \sum_{i=1}^{m} g_i(\mathbf{x})\sigma_i(\mathbf{x}). \tag{1}$$

For fixed $r \in \mathbb{N}$, one may then define a *lower bound* $\mathrm{lb}(f)_r \leq f_{\min}$ on the minimum of f over $\mathbf{X}$ by setting:

$$\mathrm{lb}(f)_r := \sup \left\{ \lambda : f(\mathbf{x}) - \lambda = \sigma_0(\mathbf{x}) + \sum_{i=1}^{m} g_i(\mathbf{x})\sigma_i(\mathbf{x}),\ \sigma_i \in \Sigma[\mathbf{x}],\ \deg(\sigma_i) \leq r \right\}.$$

The point is that while checking nonnegativity of a polynomial is hard, the parameter $\mathrm{lb}(f)_r$ may be computed by solving a *semidefinite program* of size polynomial in the number of variables n. For fixed r, and under some minor assumptions, this can be done efficiently, yielding a *tractable* bound on $f_{\min}$.

Contrary to the case of *global* nonnegativity, the classical *Positivstellensätze* of Putinar [**Put93**] and Schmüdgen [**Sch91**] show that any polynomial f *positive* on $\mathbf{X}$ has a representation of the form (1), as long as $\mathbf{X}$ satisfies a (minor) compactness condition. To be precise: Schmüdgen's result applies more generally, but requires the use of a slightly different representation (involving also products of the constraints g_i), leading to a stronger, but more computationally intensive bound $\overline{\mathrm{lb}(f)}_r$, satisfying $f_{\min} \geq \overline{\mathrm{lb}(f)}_r \geq \mathrm{lb}(f)_r$. The upshot is that under such conditions, we have *asymptotic convergence* of the hierarchies to the true minimum:

$$\lim_{r \to \infty} \mathrm{lb}(f)_r = f_{\min} \text{ and } \lim_{r \to \infty} \overline{\mathrm{lb}(f)}_r = f_{\min}.$$

A natural question is whether this asymptotic convergence may be *quantified.* From the point of view of optimization, one may see this as proving guarantees on the quality of the bounds $\mathrm{lb}(f)_r$ and $\overline{\mathrm{lb}(f)}_r$ depending on the degree r. From the point of view of real algebra, one may also see this as showing bounds on the degree of the sums of squares $\sigma_i \in \Sigma[\mathbf{x}]$ required in the decomposition of f.

As we explain in more detail in Chapter 2, we prove strong guarantees in Part 2 of this thesis when the set $\mathbf{X}$ exhibits symmetric structure. This includes the binary hypercube, the unit sphere, the standard simplex and the box $[-1, 1]^n$. There, the *polynomial kernel method* introduced in Chapter 6 allows us to express the error of Lasserre's hierarchies in terms of the behaviour of classical, univariate orthogonal polynomials. See also Chapter 1.

Measure-based sum-of-squares hierarchies

In addition to the hierarchies of *lower* bounds introduced above, Lasserre [**Las11**] also defines a hierarchy of *upper bounds* $\mathrm{ub}(f)_r$ on $f_{\min}$, which are obtained by a *structured sampling* of the feasible region $\mathbf{X}$. Apart from their inherent use in approximating $f_{\min}$, these bounds actually play a crucial role in our analysis of the hierarchies defined above. For $r \in \mathbb{N}$, the bound $\mathrm{ub}(f)_r \geq f_{\min}$ is obtained by minimizing the expectation $\int_{\mathbf{X}} f(\mathbf{x}) d\nu(\mathbf{x})$ of f on $\mathbf{X}$ over all probability measures ν of the form $d\nu(\mathbf{x}) = h(\mathbf{x}) d\mu(\mathbf{x})$, where h is a sum of squares of degree at most $2r$, and μ is a *fixed* reference measure supported on $\mathbf{X}$. These *measure-based* bounds may be computed via semidefinite programming, and asymptotic convergence to $f_{\min}$ is guaranteed when the feasible region is compact [**Las11**]. Again, this leads to natural questions about the *rate* of convergence.

As we explain in more detail in Chapter 2, the new results in Part 1 of this thesis largely settle these convergence questions (in light of several other, existing results). Indeed, we first extend a known best-possible convergence guarantee established by de Klerk & Laurent [**dKL20b, dKL20a**] for the box $[-1, 1]^n$ and hypersphere to a larger class of examples of semialgebraic sets in Chapter 3. Second, we establish a convergence rate for essentially all other semialgebraic sets which is just a log-factor away from best-possible in Chapter 4.

To prove the former result, we show in Chapter 3 that the behaviour of the measure-based bounds $\mathrm{ub}(f)_r$ depends in some sense only on the *local geometry* of the feasible region $\mathbf{X}$ near a minimizer $\mathbf{x}^*$ of f. This allows us to transport the analysis of [**dKL20b**] on $[-1, 1]^n$ to a larger class of convex bodies $\mathbf{X}$, including the unit ball and the standard simplex.

For the latter result, we construct in Chapter 4 explicit sum-of-squares densities h on $\mathbf{X}$ of degree $2r$ which approximate the *Dirac delta function* centered at a minimizer $\mathbf{x}^*$ of f. The idea is that for such h, we have:

$$\int_{\mathbf{X}} f(\mathbf{x}) h(\mathbf{x}) d\mu(\mathbf{x}) \approx f(\mathbf{x}^*) = f_{\min}.$$

That is, the density h is a good feasible solution to the program defining the upper bound $\mathrm{ub}(f)_r$. Our construction combines so-called *needle polynomials* (see Section 4.1) with *push-forward* measures (see Section 2.1). Push-forward measures were already considered by Lasserre [**Las20**] to define a more 'economical' variant of the upper bounds $\mathrm{ub}(f)_r$ (see Section 2.1). As a side result of our proof, we establish new convergence rates for these bounds as well.

Independent sets and the theta number

The *Lovász ϑ-number* is perhaps the most influential application of semidefinite programming to combinatorial optimization, providing a strong upper bound on the *independence number* $\alpha(G)$ of a graph $G = (V, E)$. The independence number is the largest cardinality of an *independent set* in G; meaning a subset $S \subseteq V$ so that no two vertices $v, w \in S$ are joined by an edge.

To each such independent set, we can associate an *incidence vector* $\mathbf{x}_S \in \mathbb{R}^V$, whose entries $(\mathbf{x}_S)_v$ are 1 if $v \in S$ and 0 otherwise. This vector, in turn, induces a positive semidefinite matrix $X = \mathbf{x}_S\mathbf{x}_S^\top/|S| \succeq 0$ which additionally satisfies $\sum_{v,w \in V} X_{v,w} = |S|$, $\mathrm{Tr}(X) = 1$ and $X_{vw} = 0$ for every edge $\{v, w\} \in E$. The idea of Lovász, now, is to compute the maximum of $\sum_{v,w \in V} X_{v,w}$ over *all* positive semidefinite matrices X with $\mathrm{Tr}(X) = 1$ and $X_{vw} = 0\, \forall \{v, w\} \in E$, which may be done efficiently using semidefinite programming. This maximum – which is known as the theta number $\vartheta(G)$ of G – is thus an upper bound on $\alpha(G)$.

Among many other possible extensions, Bachoc, Nebe, Oliveira, and Vallentin [**BNdOFV09**] extend Lovász's approach to *infinite* geometric graphs on compact metric spaces; this leads to bounds on the size of *spherical codes* and the densities of *sphere-packings*. We can view the hypersphere S^{n-1} as a graph by saying two vertices $\mathbf{x}, \mathbf{y} \in S^{n-1}$ are adjacent whenever $\mathbf{x} \cdot \mathbf{y} = t$ for some fixed $t \in \mathbb{R}$. The independence number in this case is the largest *volume* of a (measurable) set $S \subseteq S^{n-1}$ which does not contain any adjacent vertices. Such a set now has an *indicator function* $\chi_S : S^{n-1} \to \{0, 1\}$, which induces a positive *kernel* $\mathrm{K}(\mathbf{x}, \mathbf{y}) = \chi_S(\mathbf{x})\chi_S(\mathbf{y})$ on S^{n-1} with certain additional properties. The theta number in this setting is obtained by solving an (infinite-dimensional) optimization problem over all such kernels. The key point is that one may in fact restrict to kernels which are *invariant* under the symmetry of S^{n-1}. Such kernels can be classified in terms of univariate orthogonal polynomials, which leads to a more manageable formulation for ϑ. See Chapter 1.

In Chapter 10, we develop a recursive generalization of the ϑ-number for geometric *hyper*graphs on the unit sphere and on the Euclidean space, obtaining bounds on the independence number of such graphs. In the above language, a set S is independent in this context if it does not contain any k-tuple of vertices $v_1, v_2, \ldots, v_k$ which are pairwise adjacent for some fixed $k \in \mathbb{N}$ (thus $k = 2$ corresponds to the regular independence number). We call such a tuple a *k-simplex*. By exploiting symmetry, we find analytical expressions for our bounds in terms of orthogonal polynomials. An analysis of the asymptotic behaviour of these polynomials then allows us to reprove a result in *Euclidean Ramsey theory*; namely that *k-simplices* are *exponentially Ramsey*. Furthermore, we improve upon the previously known bounds for the base of the exponential.

Organization

The thesis is organized as follows. In Chapter 1, we introduce orthogonal polynomials and polynomial kernels, which form the foundation of our proof techniques throughout the rest of the thesis. In Chapter 2, we introduce the sum-of-squares hierarchies in more detail, and give an overview of known and new results on their convergence rates.

Part 1. In Chapters 3 and 4, we prove convergence results for the measure-based bounds in several different settings, covering the main results of [**SL20**] and [**SL21a**], respectively. In Chapter 5, we discuss some computational aspects of the measure-based bounds.

Part 2. In Chapter 6, we introduce the *polynomial kernel method*, which may be used to obtain convergence guarantees for sum-of-squares hierarchies on structured feasible regions. We then apply this technique in Chapters 7, 8, and 9, corresponding to the binary cube [**SL21b**], the unit ball and the standard simplex [**Slo22**], and the unit box [**LS21**], respectively.

Part 3. Finally, in Chapter 10, we introduce our new recursive theta number for geometric hypergraphs, which we apply to the unit sphere and the Euclidean space. This covers the main results of [**CSdOFSV21**]. We also include a small part of the work [**CSdOFSV22**].

Publications

This thesis is based on the following publications and preprints.

[**CSdOFSV22**] D. Castro-Silva, F.M. de Oliveira Filho, L. Slot, and F. Vallentin. A recursive theta body for hypergraphs, 2022. `arxiv:2206.03929`

[**CSdOFSV21**] D. Castro-Silva, F.M. de Oliveira Filho, L. Slot, and F. Vallentin. A recursive Lovász theta number for simplex-avoiding sets. *To appear in: Proc. Am. Math. Soc.*, 2021. `arxiv:2106.09360`

[**LS21**] M. Laurent and L. Slot. An effective version of Schmüdgen's Positivstellensatz for the hypercube. *Submitted to: Opt. Lett.*, 2021. `arXiv:2109.09528`

[**Slo22**] L. Slot. Sum-of-squares hierarchies for polynomial optimization and the Christoffel-Darboux kernel. *To appear in: SIAM J. Opt.*, 2022. `arxiv:2111.04610`

[**SL21b**] L. Slot and M. Laurent. Sum-of-squares hierarchies for binary polynomial optimization. In M. Singh and D.P. Williamson, editors, *Integer Programming and Combinatorial Optimization*, pages 43–57, Cham, 2021. Springer International Publishing. *Extended version in Math. Program.*, 2022. `https://doi.org/10.1007/s10107-021-01745-9`

[**SL21a**] L. Slot and M. Laurent. Near optimal analysis of Lasserre's univariate measure-based bounds for multivariate polynomial optimization. *Math. Program.*, 188:443–460, 2021

[**SL20**] L. Slot and M. Laurent. Improved convergence analysis of Lasserre's measure-based upper bounds for polynomial minimization on compact sets. *Math. Program.*, 193:831–871, 2020

CHAPTER 1

Orthogonal polynomials and kernel operators

We learn by rearranging what we know.

Ludwig Wittgenstein

In this chapter, we discuss some basic facts on orthogonal polynomials and polynomial kernels. The goal is to give the reader a high-level introduction to these concepts, and how they will play a role in the rest of this thesis. We prove or provide precise references for certain key results as they appear in future chapters.

1.1. Orthogonal polynomials of one variable

We begin with some facts on orthogonal polynomials of a single variable, which we shall need later. We also introduce a few special, well-known examples of such polynomials which feature in the rest of the thesis. For a comprehensive reference, see the book of Szegő [**Sze75**].

1.1.1. Basic definitions. Let μ be a finite Borel measure supported on an interval $I \subseteq \mathbb{R}$. Often, $I = [-1, 1]$ and μ will be of the form $d\mu(x) = w(x)dx$ for some continuous weight function w. We then have an inner product on the space of univariate polynomials $\mathbb{R}[x]$ by:

$$\langle f, g \rangle := \int_I f(x)g(x)d\mu(x). \tag{1.1}$$

Up to scaling, there exists a unique basis $\{P_k : k \in \mathbb{N}\}$ of $\mathbb{R}[x]$ satisfying:

$$\langle P_i, P_j \rangle = \int_I P_i(x)P_j(x)d\mu(x) = 0 \quad (i \neq j)$$

and for which $\deg(P_k) = k$ for all $k \in \mathbb{N}$. We call these P_k the *orthogonal polynomials* for the measure μ. There are several ways to normalize them. For instance, we will write:

$$\widetilde{P}_k(x) = P_k(x)/\sqrt{\langle P_k, P_k \rangle}$$

for the normalization satisfying $\langle \widetilde{P}_k, \widetilde{P}_k \rangle = 1$. We will write:

$$\overline{P}_k(x) = P_k(x)/\max_{x \in I} |P_k(x)|$$

for the normalization whose sup-norm on I is equal to 1. Both normalizations are unique up to sign.

1.1.2. The three-term reccurence relation. A useful property of orthogonal polynomials is that they satisfy a so-called *three-term recurrence relation.*

PROPOSITION 1.1 (Three-term reccurence). *Let $\{P_k : k \in \mathbb{N}\}$ be the orthogonal polynomials w.r.t. to some measure μ. Then for each $k \geq 0$, there exist constants $a_k, b_k, c_k \in \mathbb{R}$ such that:*

$$xP_k(x) = a_k P_{k+1}(x) + b_k P_k(x) + c_k P_{k-1}(x). \tag{1.2}$$

Here, we have set $P_{-1} := 0$.

The renormalizations $\widetilde{P}_k$ and $\overline{P}_k$ satisfy (1.2) as well (but with different constants). We note that for the orthonormal polynomials $\widetilde{P}_k$, the recurrence coefficients in (1.2) satisfy $c_k = a_{k-1}$ for all $k \geq 1$.

1.1.3. Extremal roots. For $r \in \mathbb{N}$, consider the matrix $J \subseteq \mathbb{R}^{(r+1)\times(r+1)}$ given by:

$$J_{i,j} = \langle x\widetilde{P}_i, \widetilde{P}_j \rangle = \int_I x\widetilde{P}_i(x)\widetilde{P}_j(x)d\mu(x) \quad (0 \leq i, j \leq r). \tag{1.3}$$

This matrix – and its eigenvalues in particular – are closely related to Lasserre's measure-based bounds (2.6), see Section 2.3. A first consequence of (1.2) is that we are able to express these eigenvalues in terms of the *roots* of the polynomials P_k. These roots are known to be real, distinct, and they lie within the interval I. As we see below, more is known about their (asymptotic) behaviour in special cases.

PROPOSITION 1.2 (see [**dKL20b**]). *Let $\{\widetilde{P}_k : k \in \mathbb{N}\}$ be orthonormal polynomials with three-term reccurence relation* (1.2)*, and let $J \subseteq \mathbb{R}^{(r+1)\times(r+1)}$ be as in* (1.3)*. Then the smallest eigenvalue $\lambda_{\min}(J)$ of J is given by:*

$$\lambda_{\min}(J) = \xi_{r+1},$$

where $\xi_{r+1} \in \mathbb{R}$ is the least root of P_{r+1}.

PROOF. We follow the proof given in [**dKL20b**]. Using the three-term reccurence (1.2), we find that:

$$\begin{aligned} J_{i,j} = \langle x\widetilde{P}_i, \widetilde{P}_j \rangle &= \langle a_i\widetilde{P}_{i+1} + b_i\widetilde{P}_i + c_i\widetilde{P}_{i-1}, \widetilde{P}_j \rangle \\ &= \begin{cases} a_i & \text{if } j = i+1, \\ b_i & \text{if } j = i, \\ c_i & \text{if } j = i-1, \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

The matrix J is thus tridiagonal, of the form:

$$J = \begin{pmatrix} b_0 & a_0 & 0 & \dots & 0 \\ c_1 & b_1 & a_1 & \dots & 0 \\ 0 & \ddots & \ddots & \ddots & \vdots \\ \vdots & \vdots & c_{r-1} & b_{r-1} & a_{r-1} \\ 0 & 0 & \dots & c_r & b_r \end{pmatrix}.$$

As we noted before, we have $c_k = a_{k-1}$, meaning J is symmetric. It follows (see [**Sze75**]) that:

$$\det(xI - J)\widetilde{P}_0 = \left(\prod_{j=0}^{r} a_j\right)\widetilde{P}_{r+1}(x),$$

which implies that the eigenvalues of J are precisely the roots of $\widetilde{P}_{r+1}$. In particular, $\lambda_{\min}(J)$ is the smallest root of $\widetilde{P}_{r+1}$. □

1.1.4. Jacobi polynomials. For parameters $\alpha, \beta > -1$, let $w_{\alpha,\beta}(x) := (1-x)^\alpha(1+x)^\beta$ be the *Jacobi weight function*. For $k \in \mathbb{N}$, we write $\mathcal{J}_k^{(\alpha,\beta)}$ for the *Jacobi polynomial* of degree k, defined by the orthogonality relation:

$$\int_{-1}^{1} \mathcal{J}_i^{(\alpha,\beta)}(x)\mathcal{J}_j^{(\alpha,\beta)}(x)w_{\alpha,\beta}(x)dx = 0 \quad (i \neq j). \tag{1.4}$$

The Jacobi polynomials satisfy the symmetry relation:

$$\mathcal{J}_k^{(\alpha,\beta)}(-x) = (-1)^k \mathcal{J}_k^{(\beta,\alpha)}(x).$$

For $\alpha \geq \beta$, the maximum $\max_{-1\leq x\leq 1} |\mathcal{J}_k^{(\alpha,\beta)}(x)|$ is attained at $x = 1$, and the normalization $\overline{\mathcal{J}}_k^{(\alpha,\beta)}$ is thus defined by setting $\overline{\mathcal{J}}_k^{(\alpha,\beta)}(1) = 1$. There are several bounds known for the roots of the Jacobi polynomials, which permit to show the following.

PROPOSITION 1.3 ([**DJ12, DN10**]. See also [**dKL20b**]). *Let $\alpha, \beta > -1$. For $k \in \mathbb{N}$, let $\mathcal{J}_k^{(\alpha,\beta)}(x)$ be the Jacobi polynomial defined in* (1.4) *and write $\xi_k \in [-1, 1]$ for its smallest root. Then we have:*

$$\xi_k = -1 + \Theta(1/k^2).$$

1.1.5. Gegenbauer polynomials. In the special case that $\alpha = \beta > -1$, the polynomials:

$$\mathcal{G}_k^{(\alpha)}(x) := \mathcal{J}_k^{(\alpha,\alpha)}(x)$$

are known as *Gegenbauer polynomials* (or *ultraspherical* polynomials). They are thus the orthogonal polynomials for the measure $d\mu(x) = (1-x^2)^\alpha dx$ on $[-1, 1]$. As we see below in Section 1.2, the Gegenbauer polynomials are connected to the unit sphere S^{n-1}; they form the basis of an analysis of sum-of-squares hierarchies on S^{n-1} in [**dKL20a**] and [**FF21**], see also Section 2.3.

We shall use them in a similar way in Chapter 8, where we analyze sum-of-squares hierarchies on the unit ball and the simplex. In Chapter 10, we use them to define bounds on the size of independent sets in S^{n-1}.

1.1.6. Chebyshev polynomials. Specializing even further to the case $\alpha = \beta = -1/2$, we get the *Chebyshev polynomials*:

$$\mathcal{C}_k(x) := \mathcal{J}_k^{(-\frac{1}{2},-\frac{1}{2})}(x).$$

Up to normalization, they may also be defined by the relation:

$$\mathcal{C}_k(x) = \cos(k \arccos x) \quad (-1 \leq x \leq 1).$$

They then satisfy a three-term recurrence with particularly simple constants:

$$x\mathcal{C}_k(x) = \frac{1}{2}\mathcal{C}_{k+1}(x) + \frac{1}{2}\mathcal{C}_{k-1}(x).$$

The Chebyshev polynomials appear often in mathematical optimization and they satisfy several useful extremal properties, see for instance Theorem 7.29. They are used to define so-called *needle polynomials* (see Section 4.1), which approximate well the Dirac delta function on $[-1, 1]$, as well as the well-known *Jackson kernel*, which is used in functional approximation. These constructions play a central role in Chapter 4 and Chapter 9, respectively.

1.1.7. Krawtchouk polynomials. Finally, we consider for $n \in \mathbb{N}$ the *Krawtchouk polynomials* $\mathcal{K}_k^{(n)}(x)$, $0 \leq k \leq n$, which are given by:

$$\mathcal{K}_k^{(n)}(x) := \sum_{i=0}^{k} (-1)^i \binom{x}{i}\binom{n-x}{k-i} \quad (0 \leq k \leq n).$$

They are the orthogonal polynomials with respect to the *discrete* measure $\mu = \frac{1}{2^n}\sum_{x=0}^{n}\binom{n}{x}\delta_x$ on $[0, n]$, where δ_x is the Dirac measure centered at x. Note that this is a slight departure from the previous setting, as μ now has finite support. As a consequence, the inner product (1.1) is defined only on the space $\mathbb{R}[\mathbf{x}]_n$ of polynomials of degree at most n, but all other mentioned results carry over. The maximum $\max_{0 \leq x \leq n} |\mathcal{K}_k^{(n)}(x)|$ is attained at $x = 0$, and the normalization $\overline{\mathcal{K}}_k^{(n)}$ is thus defined by setting $\overline{\mathcal{K}}_k^{(n)}(0) = 1$.

The Krawtchouk polynomials are related to the binary hypercube $\{0, 1\}^n$, see Section 1.2. Indeed, they will feature prominently in Chapter 7, where we analyze sum-of-squares hierarchies on $\{0, 1\}^n$. They also appear in Chapter 10 to define bounds on the size of independent sets on $\{0, 1\}^n$.

The asymptotic behaviour of the least root of $\mathcal{K}_r^{(n)}(x)$ is studied by Levenshtein [**Lev98**] in the regime $r/n \to t \in [0, 1]$, see also Theorem 7.2.

1.2. Polynomial kernels and kernel operators

Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact set, and let μ be a finite positive Borel measure supported on $\mathbf{X}$. Consider the space of polynomials $\mathcal{P}(\mathbf{X})$ restricted to $\mathbf{X}$, which is given by:

$$\mathcal{P}(\mathbf{X}) = \mathbb{R}[\mathbf{x}]/(p : p(\mathbf{x}) = 0 \quad \forall \mathbf{x} \in \mathbf{X}).$$

It is thus the quotient of the polynomial ring modulo the *vanishing ideal* of $\mathbf{X}$, which is the ideal generated by the polynomials that vanish on $\mathbf{X}$. The degree of an element $p \in \mathcal{P}(\mathbf{X})$ is the smallest degree of a polynomial $h \in \mathbb{R}[\mathbf{x}]$ with $p(\mathbf{x}) = h(\mathbf{x})$ for all $\mathbf{x} \in \mathbf{X}$. For an integer d, we write $\mathcal{P}(\mathbf{X})_d$ for the polynomials in $\mathcal{P}(\mathbf{X})$ of degree at most d. We have an inner product on $\mathcal{P}(\mathbf{X})$ given by:

$$\langle f, g \rangle_\mu = \int_{\mathbf{X}} f(\mathbf{x}) g(\mathbf{x}) d\mu(\mathbf{x}).$$

Now let $\mathrm{K} : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ be a *polynomial kernel*, meaning that $\mathrm{K}(\mathbf{x}, \mathbf{y})$ is a polynomial in $\mathbf{x}$ and $\mathbf{y}$. Using the inner product $\langle \cdot, \cdot \rangle_\mu$, we can associate a linear operator $\mathbf{K} : \mathcal{P}(\mathbf{X}) \to \mathcal{P}(\mathbf{X})$ to the kernel K via:

$$\mathbf{K}p(\mathbf{x}) := \langle \mathrm{K}(\mathbf{x}, \cdot), p \rangle_\mu = \int_{\mathbf{X}} \mathrm{K}(\mathbf{x}, \mathbf{y}) p(\mathbf{y}) d\mu(\mathbf{y}). \tag{1.5}$$

1.2.1. The Christoffel-Darboux kernel. The space $\mathcal{P}(\mathbf{X})$ has an orthonormal basis w.r.t. the inner product $\langle \cdot, \cdot \rangle_\mu$ given by polynomials $\{P_\alpha : \alpha \in A\}$ of degree $|\alpha|$ satisfying:

$$\langle P_\alpha, P_\beta \rangle = \int_{\mathbf{X}} P_\alpha(\mathbf{x}) P_\beta(\mathbf{x}) d\mu(\mathbf{x}) = \delta_{\alpha\beta} \quad (\alpha, \beta \in A). \tag{1.6}$$

Here, the set $A \subseteq \mathbb{N}^n$ depends on $\mathbf{X}$. For instance, if $\mathbf{X}$ is full-dimensional, we simply have $A = \mathbb{N}^n$ (as $\mathcal{P}(\mathbf{X}) = \mathbb{R}[\mathbf{x}]$). However, if for example $\mathbf{X} = \{0,1\}^n$ is the binary hypercube, we have $\mathcal{P}(\mathbf{X}) = \mathrm{span}\{x^\alpha : \alpha_i \in \{0,1\}\}$, meaning $A = \{0,1\}^n \subset \mathbb{N}^n$. In what follows, the set A will be clear from the context, and so we will not always denote it explicitely. We also assume throughout that the P_α are chosen so that the space of polynomials $\mathcal{P}(\mathbf{X})_d$ of polynomials of degree at most d is spanned by the polynomials P_α with $|\alpha| \leq d$ for each $d \in \mathbb{N}$.

Using such an orthonormal basis, we are able to construct the so-called *Christoffel-Darboux kernel* CD_r of degree $r \in \mathbb{N}$, which is defined as:

$$\mathrm{CD}_r(\mathbf{x}, \mathbf{y}) = \sum_{\alpha \in A : |\alpha| \leq r} P_\alpha(\mathbf{x}) P_\alpha(\mathbf{y}). \tag{1.7}$$

The point is that the operator $\mathbf{CD}_r$ associated to this kernel via (1.5) *reproduces* the space of polynomials $\mathcal{P}(\mathbf{X})_r$ on $\mathbf{X}$ of degree at most r, meaning that:

$$\mathbf{CD}_r p(\mathbf{x}) = p(\mathbf{x}) \tag{1.8}$$

for all $p \in \mathcal{P}(\mathbf{X})$ of degree at most r. Indeed, such a p may be written as:

$$p(\mathbf{x}) = \sum_{|\alpha| \leq r} p_\alpha P_\alpha(\mathbf{x})$$

for certain $p_\alpha \in \mathbb{R}$, whence:

$$\begin{aligned}\mathbf{CD}_r p(\mathbf{x}) &= \sum_{|\alpha| \leq r, |\beta| \leq r} p_\alpha \int_{\mathbf{X}} P_\beta(\mathbf{x}) P_\beta(\mathbf{y}) P_\alpha(\mathbf{y}) d\mu(\mathbf{y}) \\ &= \sum_{|\alpha| \leq r} p_\alpha P_\alpha(\mathbf{x}) = p(\mathbf{x})\end{aligned}$$

using the orthogonality relations (1.6). Note that for $\mathbf{x}, \mathbf{y} \in \mathbf{X}$, the value of $\mathrm{CD}_r(\mathbf{x}, \mathbf{y})$ does not depend on the choice of basis $\{P_\alpha\}$. Indeed, the subspace generated by the P_α of degree at most r does not depend on the choice of basis (see also (1.10) below). Therefore, for each $\mathbf{y} \in \mathbf{X}$, the polynomial $\mathbf{x} \mapsto \mathrm{CD}_r(\mathbf{x}, \mathbf{y})$ of degree r in $\mathcal{P}(\mathbf{X})$ is fixed by the reproducing property (1.8).

The Christoffel-Darboux kernel has many applications in optimization, and it will form the basis of a technique we discuss in Chapter 6 to analyze sum-of-squares hierarchies for polynomial optimization. More specifically, we will make use there of a class of kernels defined in the spirit of (1.7) as:

$$\mathrm{CD}_r(\mathbf{x}, \mathbf{y}; \lambda) := \sum_{|\alpha| \leq r} \lambda_\alpha P_\alpha(\mathbf{x}) P_\alpha(\mathbf{y}) \quad (\lambda_\alpha \in \mathbb{R}), \tag{1.9}$$

which we call *perturbed* Christoffel-Darboux kernels. Using again the relations (1.6), the operator $\mathbf{CD}_r(\lambda)$ associated to this kernel is diagonal w.r.t. the basis $\{P_\alpha\}$, and its eigenvalues are given by the coefficients λ_α.

1.2.2. Summation formulas. For certain special, structured sets $\mathbf{X}$ and measures μ, the Christoffel-Darboux kernel (1.7) admits a simple, closed form expression in terms of *univariate* orthogonal polynomials. These expressions will be of great help in Part 2 and Part 3 of the thesis, where we consider for $\mathbf{X}$ the binary cube $\{0,1\}^n$, the unit sphere S^{n-1}, the unit ball B^n and the standard simplex Δ^n.

The basic idea is to consider the subspaces $H_k \subseteq \mathcal{P}(\mathbf{X})$ given by:

$$\begin{aligned}H_k &= \mathrm{span}\{P_\alpha : |\alpha| = k\} \\ &= \{p \in \mathcal{P}(\mathbf{X}) : \langle p, q \rangle_\mu = 0 \text{ for all } q \in \mathcal{P}(\mathbf{X})_{k-1}\}.\end{aligned} \tag{1.10}$$

Note that the H_k depend on $\mathbf{X}$ and μ, but not on the choice of basis $\{P_\alpha\}$. We now have the orthogonal decomposition:

$$\mathcal{P}(\mathbf{X}) = \bigoplus_{k=0}^{\infty} H_k. \tag{1.11}$$

Accordingly, we can write any $p \in \mathcal{P}(\mathbf{X})_d$ of degree d as:

$$p(\mathbf{x}) = p_0(\mathbf{x}) + p_1(\mathbf{x}) + \ldots + p_d(\mathbf{x}) \quad (p_k \in H_k), \tag{1.12}$$

and similarly, the Christoffel-Darboux kernel (1.7) can be written as:

$$\mathrm{CD}_r(\mathbf{x}, \mathbf{y}) = \sum_{k=0}^{r} \mathrm{CD}^{(k)}(\mathbf{x}, \mathbf{y}), \text{ where } \mathrm{CD}^{(k)}(\mathbf{x}, \mathbf{y}) := \sum_{|\alpha|=k} P_\alpha(\mathbf{x}) P_\alpha(\mathbf{y}). \quad (1.13)$$

The operator $\mathbf{CD}^{(k)}$ associated to $\mathrm{CD}^{(k)}$ then satisfies:

$$\mathbf{CD}^{(k)} p(\mathbf{x}) = p_k(\mathbf{x}). \quad (1.14)$$

In the special cases mentioned above, the kernels $\mathrm{CD}^{(k)}(\mathbf{x}, \mathbf{y})$ may be expressed in terms of the orthogonal polynomials of Section 1.1. For illustration, we briefly cover the binary cube and the unit sphere below, which are the most classical examples. See for instance [**Val08**]. The unit ball and standard simplex are treated in Chapter 8.

The unit sphere. Consider the unit sphere $\mathbf{X} = S^{n-1} \subseteq \mathbb{R}^n$, equipped with the uniform surface measure μ. The space $\mathcal{P}(S^{n-1})$ of polynomials on S^{n-1} is given by:

$$\mathcal{P}(S^{n-1}) = \mathbb{R}[\mathbf{x}]/(1 - \|\mathbf{x}\|^2).$$

The subspaces H_k in this case are given by:

$$H_k = \mathrm{Harm}_k := \{p : \text{ is homogeneous of degree } k \text{ and harmonic}\}.$$

A polynomial p is harmonic if it is in the kernel of the Laplace operator, i.e., if $\nabla^2 p = 0$. An element of Harm_k is also called a *spherical harmonic* (of degree k). Choosing any orthonormal basis $\{s_{k,j}\}$ of Harm_k, and for the right normalization of the Gegenbauer polynomials $\mathcal{G}_k^{(\alpha)}$, we have the summation formula:

$$\mathrm{CD}^{(k)}(\mathbf{x}, \mathbf{y}) = \sum_j s_{k,j}(\mathbf{x}) s_{k,j}(\mathbf{y}) = \mathcal{G}_k^{(\frac{n-3}{2})}(\mathbf{x} \cdot \mathbf{y}). \quad (1.15)$$

What this means is that the Christoffel-Darboux kernel (1.13) is given by:

$$\mathrm{CD}_r(\mathbf{x}, \mathbf{y}) = \sum_{k=0}^{r} \mathcal{G}_k^{(\frac{n-3}{2})}(\mathbf{x} \cdot \mathbf{y}).$$

If we choose coeffcients $\lambda_\alpha = \lambda_{|\alpha|}$ depending only on $|\alpha|$, the perturbed kernel (1.9) is then given by:

$$\mathrm{CD}_r(\mathbf{x}, \mathbf{y}; \lambda) = \sum_{k=0}^{r} \lambda_k \mathcal{G}_k^{(\frac{n-3}{2})}(\mathbf{x} \cdot \mathbf{y}).$$

Another way of looking at this is as follows. Let $u \in \mathbb{R}[x]$ be a univariate polynomial of degree $r \geq d$, with the following expression in the basis of Gegenbauer polynomials:

$$u(x) = \sum_{k=0}^{r} \lambda_k \mathcal{G}_k^{(\frac{n-3}{2})}(x),$$

and consider the kernel $\mathrm{K}(\mathbf{x},\mathbf{y}) = u(\mathbf{x}\cdot\mathbf{y})$. Then, we have:

$$\mathrm{K}(\mathbf{x},\mathbf{y}) = \sum_{k=0}^{r} \lambda_k \mathcal{G}_k^{(\frac{n-3}{2})}(\mathbf{x}\cdot\mathbf{y}) = \mathrm{CD}_r(\mathbf{x},\mathbf{y};\lambda) \tag{1.16}$$

and thus, by (1.14), we get the *Funk-Hecke formula*:

$$\mathbf{K}p(\mathbf{x}) = \int_{\mathbf{X}} u(\mathbf{x}\cdot\mathbf{y})p(\mathbf{y})d\mu(\mathbf{y}) = \sum_{k=0}^{d} \lambda_k p_k(\mathbf{x}),$$

where $p(\mathbf{x}) = \sum_{k=0}^{d} p_k(\mathbf{x}),\ p_k \in \mathrm{Harm}_k$ is as in (1.12).

The binary cube. Consider the binary cube $\mathbf{X} = \{0,1\}^n \subseteq \mathbb{R}^n$, equipped with the uniform probability measure μ. The space $\mathcal{P}(\{0,1\}^n)$ is spanned by *multilinear* polynomials, i.e., we have:

$$\mathcal{P}(\{0,1\}^n) = \mathbb{R}[\mathbf{x}]/(\mathbf{x}_i - \mathbf{x}_i^2 \ :\ 1 \leq i \leq n) = \mathrm{span}\{\mathbf{x}^\alpha : \alpha \in \{0,1\}^n\}.$$

The decomposition (1.11) in this case is given in terms of the *characters*:

$$\chi_a(\mathbf{x}) := (-1)^{\mathbf{x}\cdot a} = \prod_{i:a_i=1} (1 - 2\mathbf{x}_i) \quad (a \in \{0,1\}^n),$$

which are polynomials of degree $|a|$ (on the binary cube). They form an orthonormal basis of $\mathcal{P}(\{0,1\}^n)$, and the spaces H_k of (1.10) are given by:

$$H_k = \mathrm{span}\{\chi_a : |a| = k\} \quad (0 \leq k \leq n).$$

Similar to the spherical harmonics, we have a summation formula for the characters, now in terms of Krawtchouk polynomials:

$$\sum_{|a|=k} \chi_a(\mathbf{x})\chi_a(\mathbf{y}) = \mathcal{K}_k^{(n)}(d_{\mathrm{ham}}(\mathbf{x},\mathbf{y})), \tag{1.17}$$

where $d_{\mathrm{ham}}(\mathbf{x},\mathbf{y}) = |\{i : \mathbf{x}_i \neq \mathbf{y}_i\}|$ is the *Hamming distance* between $\mathbf{x}$ and $\mathbf{y}$. Now let $u \in \mathbb{R}[x]$ be a univariate polynomial of degree $n \geq r \geq d$, expressed in the basis of Krawtchouk polynomials as:

$$u(x) = \sum_{k=0}^{r} \lambda_k \mathcal{K}_k^{(n)}(x),$$

and consider the kernel $\mathrm{K}(\mathbf{x},\mathbf{y}) = u(d_{\mathrm{ham}}(\mathbf{x},\mathbf{y}))$. Then we again get the Funk-Hecke formula:

$$\mathbf{K}p(\mathbf{x}) = \frac{1}{2^n} \sum_{\mathbf{y}\in\{0,1\}^n} u(d_{\mathrm{ham}}(\mathbf{x},\mathbf{y}))p(\mathbf{y}) = \sum_{k=0}^{d} \lambda_k p_k(\mathbf{x})$$

for $p(\mathbf{x}) = \sum_{k=0}^{d} p_k(\mathbf{x}),\ p_k \in H_k$ decomposed as in (1.12).

1.2.3. Symmetry and invariant kernels. An aspect we have so far ignored is the role of *symmetry*. Both the unit sphere and binary cube are highly symmetrical sets, and their symmetry is connected to the representations of the Christoffel-Darboux kernel given above.

The unit sphere. For S^{n-1}, we have the regular $O(n)$-action, which induces an action on $\mathcal{P}(S^{n-1})$ via $Tf(\mathbf{x}) = f(T^{-1}\mathbf{x})$. It turns out that the spaces Harm_k are precisely the *invariant* and *irreducible* orthogonal components of $\mathcal{P}(S^{n-1})$ under this action. What this means is that each of the spaces Harm_k is closed onder the $O(n)$-action, and that it contains no proper closed subspaces. The inner product $\mathbf{x} \cdot \mathbf{y}$ is *invariant* under the $O(n)$-action, meaning that $\mathbf{x} \cdot \mathbf{y} = T(\mathbf{x}) \cdot T(\mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in S^n$ and $T \in O(n)$. Therefore, kernels (1.16) of the form $\mathrm{K}(\mathbf{x}, \mathbf{y}) = u(\mathbf{x} \cdot \mathbf{y})$ satisfy:

$$\mathrm{K}(\mathbf{x}, \mathbf{y}) = \mathrm{K}(T\mathbf{x}, T\mathbf{y}) \qquad (\mathbf{x}, \mathbf{y} \in S^{n-1},\ T \in O(n)).$$

Such kernels are called $O(n)$-invariant. In fact, *all* invariant polynomial kernels on S^{n-1} are of this form, see [**Val08**].

The binary cube. For the binary cube $\{0,1\}^n$, we have an action of $(\mathbb{Z}/2\mathbb{Z})^n \rtimes S_n$ which is generated by coordinate permutations $\mathbf{x} \mapsto \sigma(\mathbf{x}) = (\mathbf{x}_{\sigma(1)}, \ldots, \mathbf{x}_{\sigma(n)})$, $\sigma \in S_n$, and 'bit-flips' $\mathbf{x} \mapsto a \oplus \mathbf{x} = a + \mathbf{x} \mod 2$, $a \in (\mathbb{Z}/2\mathbb{Z})^n$. The spaces H_k spanned by the characters of exact degree k are invariant and irreducible under this action. Furthermore, the Hamming distance $d_{\mathrm{ham}}(\mathbf{x}, \mathbf{y})$ is invariant, and the invariant kernels on $\{0,1\}^n$ are all of the form $\mathrm{K}(\mathbf{x}, \mathbf{y}) = u(d_{\mathrm{ham}}(\mathbf{x}, \mathbf{y}))$, see [**Val08**].

1.2.4. Preview of applications. We will use kernels and their associated operators in two ways. First, as we explain in Chapter 6, one can deduce guarantees on the quality of sum-of-squares hierarchies by constructing kernel operators having certain 'nice' properties. In the special cases where one has a representation of the type (1.16) for $\mathrm{CD}_r(\mathbf{x}, \mathbf{y}; \lambda)$ in terms of a univariate polynomial $u \in \mathbb{R}[x]$, one may express these properties in terms of the coefficients λ_k of this polynomial u in the appropriate basis of orthogonal polynomials. This effectively reduces the problem from a multivariate to a univariate setting.

Second, as we see in Chapter 10, one may define bounds on the size of independent sets in (hyper)graphs using kernels of *positive type*. On the unit sphere and binary cube, invariant kernels of this type can be classified using the expression (1.16); namely, such kernels are of positive type if and only if the coefficients λ_k are nonnegative. The resulting bounds may then be studied by analyzing certain asymptotic properties of the corresponding orthogonal polynomials.

CHAPTER 2

Polynomial optimization and sum-of-squares hierarchies

You know, this is, excuse me, a damn fine cup of coffee.

Dale Cooper, Twin Peaks

In this chapter, we introduce several hierarchies of semidefinite relaxations for polynomial optimization problems based on sums of squares. Their analysis will be the subject of Part 1 and Part 2 of the thesis.

A *polynomial optimization problem* asks to minimize a given n-variate polynomial $f \in \mathbb{R}[\mathbf{x}]$ over a *semialgebraic set* $\mathbf{X} \subseteq \mathbb{R}^n$, itself defined by polynomials $g_1, g_2, \ldots, g_m \in \mathbb{R}[\mathbf{x}]$ as:

$$\mathbf{X} = \{\mathbf{x} \in \mathbb{R}^n : g_j(\mathbf{x}) \geq 0 \quad (1 \leq j \leq m)\}. \tag{2.1}$$

That is, it asks to compute the global minimum:

$$f_{\min} := \inf_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x}). \tag{2.2}$$

Polynomial optimization is generally hard and non-convex. Many classical combinatorial problems including MaxCut and StableSet may be formulated as a polynomial optimization problem, already for simple feasible regions $\mathbf{X}$, such as the unit sphere, the unit ball, the binary hypercube or the standard simplex. For instance, one may compute the stability number of a graph $G = (V, E)$ as:

$$\alpha(G) = \max_{\mathbf{x} \in \{0,1\}^V} \sum_{i \in V} \mathbf{x}_i - \sum_{\{i,j\} \in E} \mathbf{x}_i \mathbf{x}_j.$$

This is equivalent to (see [**PH13**]):

$$\alpha(G) = \max_{\mathbf{x} \in [0,1]^V} \sum_{i \in V} \mathbf{x}_i - \sum_{\{i,j\} \in E} \mathbf{x}_i \mathbf{x}_j.$$

Hence, $\alpha(G)$ may be formulated as a polynomial optimization problem over both the binary and continuous hypercube. The formulation (2.2) in fact naturally includes (binary) linear programming and quadratic programming.

Alternatively, we may formulate $\alpha(G)$ as (see [**MS65**]):

$$\frac{1}{\alpha(G)} = \min_{\mathbf{x}\in\mathbb{R}^n} \left\{ \sum_{i\in V} \mathbf{x}_i^2 + 2 \sum_{\{i,j\}\in E} \mathbf{x}_i\mathbf{x}_j : \mathbf{x} \geq 0,\ \sum_{i\in V} \mathbf{x}_i = 1 \right\},$$

or similarly as:

$$\frac{1}{\alpha(G)} = \min_{\mathbf{x}\in S^{n-1}} \sum_{i\in V} \mathbf{x}_i^4 + 2 \sum_{\{i,j\}\in E} \mathbf{x}_i^2\mathbf{x}_j^2.$$

Thus, as polynomial optimization problems over the unit sphere and the standard simplex.

The general intractability of (2.2) motivates the search for efficient *bounds* on $f_{\min}$. In this chapter, we consider bounds that are based on *sums of squares* of polynomials. A polynomial $\sigma \in \mathbb{R}[\mathbf{x}]$ is a sum of squares if there exist polynomials $p_1, p_2, \ldots, p_\ell$ such that:

$$\sigma(\mathbf{x}) = p_1(\mathbf{x})^2 + p_2(\mathbf{x})^2 + \ldots + p_\ell(\mathbf{x})^2.$$

We write $\Sigma[\mathbf{x}] \subseteq \mathbb{R}[\mathbf{x}]$ for the set of all sum-of-squares polynomials. We denote by $\Sigma[\mathbf{x}]_r$ the restriction of $\Sigma[\mathbf{x}]$ to polynomials of degree at most r. Importantly, if $\sigma \in \Sigma[\mathbf{x}]$ is a sum of squares, it is *globally nonnegative*, i.e., $\sigma(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$. Lasserre [**Las01**] and Parillo [**Par00**] use sums of squares to define *hierarchies* of *lower bounds* on the minimum $f_{\min}$ of f. At fixed *level* r of the hierarchy, these bounds may be computed by solving a *semidefinite program* of size polynomial in the number of variables n, as we see in more detail below. Lasserre [**Las11, Las20**] also uses sums of squares to define hierarchies of *upper bounds* on $f_{\min}$. These bounds may similarly be computed using semidefinite programming.

One of the key features of these hierarchies is that they converge to the true minimum $f_{\min}$ as the level r goes to infinity (under certain mild assumptions). A natural question – which dominates a large portion of this thesis – is whether one may *quantify* this convergence, that is whether one may prove guarantees on the quality of the approximations as a function of the level r.

In the rest of this chapter, we will first introduce the sum-of-squares hierarchies in more detail. Then, we cover the basics of semidefinite programming, and show how it may be used to compute the resulting bounds. Next, we discuss existing results on their rates of convergence and some of the techniques used to establish them. Finally, we summarize the main results of the first two parts of this thesis; namely we outline new and improved convergence rates for both the upper and lower bounds in a large variety of settings.

2.1. Certificates of nonnegativity

The program (2.2) may be reformulated as finding the largest $\lambda \in \mathbb{R}$ for which the polynomial $f - \lambda$ is nonnegative on $\mathbf{X}$. That is, writing $\mathcal{P}_+(\mathbf{X}) \subseteq$

$\mathbb{R}[\mathbf{x}]$ for the cone of all polynomials that are nonnegative on $\mathbf{X}$, we have:

$$f_{\min} = \sup\{\lambda \in \mathbb{R} : f - \lambda \in \mathcal{P}_+(\mathbf{X})\}.$$

This reformulation of (2.2) establishes a connection between polynomial optimization and the problem of certifying nonnegativity of a polynomial over a semialgebraic set. Using this connection and *certificates of nonnegativity* for polynomials on compact semialgebraic sets based on sums of squares, Lasserre [**Las01**] and Parillo [**Par00**] introduce several hierarchies of bounds on $f_{\min}$.

Consider the *quadratic module* $\mathcal{Q}(\mathbf{X})$ and the *preordering* $\mathcal{T}(\mathbf{X})$ of $\mathbf{X}$, defined as:

$$\mathcal{Q}(\mathbf{X}) := \Big\{ \sum_{j=0}^{m} \sigma_j g_j : \sigma_j \in \Sigma[\mathbf{x}] \Big\} \qquad \text{(where } g_0 := 1\text{)},$$
$$\mathcal{T}(\mathbf{X}) := \Big\{ \sum_{J \subseteq [m]} \sigma_J g_J : \sigma_J \in \Sigma[\mathbf{x}] \Big\} \quad \text{(where } g_J := \prod_{j \in J} g_j\text{)}.$$

Note that strictly speaking, $\mathcal{Q}(\mathbf{X})$ and $\mathcal{T}(\mathbf{X})$ do not depend on the set $\mathbf{X}$, but rather on its description (2.1) as a semialgebraic set. We adopt this slight abuse of notation for clarity of exposition, as canonical descriptions are available for each of the sets $\mathbf{X}$ we consider. As sum-of-squares polynomials are globally nonnegative, it is clear that:

$$\Sigma[\mathbf{x}] \subseteq \mathcal{Q}(\mathbf{X}) \subseteq \mathcal{T}(\mathbf{X}) \subseteq \mathcal{P}_+(\mathbf{X}).$$

One may thus verify nonnegativity of a polynomial f over $\mathbf{X}$ by showing that f lies either in $\Sigma[\mathbf{x}]$, $\mathcal{Q}(\mathbf{X})$ or $\mathcal{T}(\mathbf{X})$.

2.1.1. Hierarchies of lower bounds. The key observation of Lasserre [**Las01**] is that membership in the *truncated* quadratic module or preordering, defined as:

$$\mathcal{Q}(\mathbf{X})_{2r} := \Big\{ \sum_{j=0}^{m} \sigma_j g_j : \sigma_j \in \Sigma[\mathbf{x}],\ \deg(\sigma_j g_j) \leq 2r \Big\},$$
$$\mathcal{T}(\mathbf{X})_{2r} := \Big\{ \sum_{J \subseteq [m]} \sigma_J g_J : \sigma_J \in \Sigma[\mathbf{x}],\ \deg(\sigma_J g_J) \leq 2r \Big\},$$

may be checked by solving a semidefinite program whose size depends on n, m and r. This leads to the following hierarchies of *lower* bounds on the global minimum $f_{\min}$ of f on $\mathbf{X}$:

$$\text{lb}(f, \mathcal{Q}(\mathbf{X}))_r := \sup\{\lambda \in \mathbb{R} : f - \lambda \in \mathcal{Q}(\mathbf{X})_{2r}\}, \tag{2.3}$$
$$\text{lb}(f, \mathcal{T}(\mathbf{X}))_r := \sup\{\lambda \in \mathbb{R} : f - \lambda \in \mathcal{T}(\mathbf{X})_{2r}\}. \tag{2.4}$$

By definition, we have $\text{lb}(f, \mathcal{Q}(\mathbf{X}))_r \leq \text{lb}(f, \mathcal{T}(\mathbf{X}))_r \leq f_{\min}$ for all $r \in \mathbb{N}$. Furthermore, the bounds converge to the global minimum $f_{\min}$ as $r \to \infty$ under mild assumptions on $\mathbf{X}$. This is a consequence of the *Positivstellensätze* of Putinar and Schmüdgen, respectively.

THEOREM 2.1 (Putinar's Positivstellensatz [**Put93**]). *Let* $\mathbf{X} \subseteq \mathbb{R}^n$ *be a semialgebraic set, and assume that* $R - \|\mathbf{x}\|^2 \in \mathcal{Q}(\mathbf{X})$ *for some* $R > 0$. *Then for any polynomial* $f \in \mathcal{P}_+(\mathbf{X})$ *and* $\eta > 0$, *we have* $f + \eta \in \mathcal{Q}(\mathbf{X})$.

THEOREM 2.2 (Schmüdgen's Positivstellensatz [**Sch91**]). *Let* $\mathbf{X} \subseteq \mathbb{R}^n$ *be a compact semialgebraic set. Then for any polynomial* $f \in \mathcal{P}_+(\mathbf{X})$ *and* $\eta > 0$, *we have* $f + \eta \in \mathcal{T}(\mathbf{X})$.

Semialgebraic sets $\mathbf{X}$ for which $R - \|\mathbf{x}\|^2$ lies in the quadratic module $\mathcal{Q}(\mathbf{X})$ for some $R > 0$ satisfy the so-called *Archimedean condition.* Note that such sets must be compact, and the requirement put on $\mathbf{X}$ in Theorem 2.1 is thus stronger than the one in Theorem 2.2. Again, we note that our notation assumes that $\mathbf{X}$ is implicitly equipped with a description (2.1).

It should also be noted that polynomials positive on $\mathbf{X}$ need not be sum-of-squares (they need not even be *globally* nonnegative). Therefore, a hierarchy of relaxations of the type (2.3) where one instead demands that $f - \lambda \in \Sigma[\mathbf{x}]_{2r}$ does not converge to $f_{\min}$ in general (the relaxation might not be feasible for *any* $r \in \mathbb{N}$). This is true even when $\mathbf{X} = \mathbb{R}^n$ (for $n \geq 2$).

EXAMPLE 2.3. *Consider the* Motzkin polynomial

$$f(\mathbf{x}_1, \mathbf{x}_2) = \mathbf{x}_1^2\mathbf{x}_2^2(\mathbf{x}_1^2 + \mathbf{x}_2^2 - 3) + 1.$$

Global nonnegativity of f *(on* $\mathbb{R}^2$*) is a consequence of the arithmetic-geometric mean inequality:*

$$\frac{1 + \mathbf{x}_1\mathbf{x}_2^4 + \mathbf{x}_1^4\mathbf{x}_2^2}{3} \geq \sqrt[3]{\mathbf{x}_1^6\mathbf{x}_2^6}.$$

However, f *is famously not a sum of squares. In fact, it is well-known that* $f + \lambda$ *is not a sum of squares for any* $\lambda \geq 0$, *see* [**Rez00**] *or* [**Lau09**]. *For any* $\lambda > 0$, *however, Putinar's and Schmüdgen's Positivstellensätze state that* $f + \lambda$ *lies in* $\mathcal{Q}([-1,1]^2)_r$ *and* $\mathcal{T}([-1,1]^2)_r$ *for* $r = r(\lambda) \in \mathbb{N}$ *large enough. In other words, the bounds* $\mathrm{lb}(f, \mathcal{Q}([-1,1]^2))_r$ *and* $\mathrm{lb}(f, \mathcal{T}([-1,1]^2))_r$ *converge to* $f_{\min} = 0$ *as* $r \to \infty$, *whereas a relaxation using only sums of squares would not be feasible at any level (and thus give no information on the minimum of* f *over* $[-1,1]^2$*).*

2.1.2. Hierarchies of upper bounds. An alternative way to reformulate problem (2.2) is as follows:

$$f_{\min} = \inf_{\nu \in \mathcal{M}(\mathbf{X})} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) d\nu(\mathbf{x}) : \int_{\mathbf{X}} d\nu(\mathbf{x}) = 1 \right\}. \tag{2.5}$$

Here $\mathcal{M}(\mathbf{X})$ denotes the set of (positive) measures supported on $\mathbf{X}$. Indeed, we see that the optimum value of (2.5) must be at least $f_{\min}$, as we are taking the expectation of f w.r.t. some probability measure on $\mathbf{X}$. On the other hand, choosing for ν the Dirac measure centered in a minimizer of f over $\mathbf{X}$ shows that the optimum value of (2.5) is at most $f_{\min}$.

The idea of Lasserre [**Las11**] now is to optimize not over the full set of measures on $\mathbf{X}$, but only over measures of the form $d\nu(\mathbf{x}) = q(\mathbf{x})d\mu(\mathbf{x})$, where

μ is a *fixed* reference measure supported on $\mathbf{X}$, and $q \in \mathbb{R}[\mathbf{x}]_{2r}$ is a polynomial known to be nonnegative on $\mathbf{X}$. Such a relaxation yields an *upper* bound on the global minimum $f_{\min}$. Based on this observation, Lasserre [**Las11**] defines for $r \in \mathbb{N}$:

$$\mathrm{ub}(f, \mathbf{X}, \mu)_r := \inf_{q \in \Sigma[\mathbf{x}]_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mu(\mathbf{x}) : \int_{\mathbf{X}} q(\mathbf{x}) d\mu(\mathbf{x}) = 1 \right\}, \tag{2.6}$$

$$\mathrm{ub}(f, \mathcal{Q}(\mathbf{X}), \mu)_r := \inf_{q \in \mathcal{Q}(\mathbf{X})_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mu(\mathbf{x}) : \int_{\mathbf{X}} q(\mathbf{x}) d\mu(\mathbf{x}) = 1 \right\}, \tag{2.7}$$

$$\mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r := \inf_{q \in \mathcal{T}(\mathbf{X})_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mu(\mathbf{x}) : \int_{\mathbf{X}} q(\mathbf{x}) d\mu(\mathbf{x}) = 1 \right\}. \tag{2.8}$$

The value of these parameters depends on the choice of reference measure μ. When this measure is clear from the context, we will sometimes leave it out of the notation, especially when μ is the Lebesgue measure restricted to $\mathbf{X}$. Note that the bound $\mathrm{ub}(f, \mathbf{X}, \mu)_r$ can be defined even if $\mathbf{X}$ is not a semialgebraic set.

As was the case for the lower bounds, each of the upper bounds can be computed by solving a semidefinite program whose size depends on n, r (and the number of inequalities m that define $\mathbf{X}$ in the case of $\mathrm{ub}(f, \mathcal{Q}(\mathbf{X}), \mu)_r$ and $\mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r$). They satisfy:

$$f_{\min} \leq \mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r \leq \mathrm{ub}(f, \mathcal{Q}(\mathbf{X}), \mu)_r \leq \mathrm{ub}(f, \mathbf{X}, \mu)_r \quad (r \in \mathbb{N}).$$

In contrast to the lower bounds, the upper bounds $\mathrm{ub}(f, \mathbf{X}, \mu)_r$ obtained by optimizing over $q \in \Sigma[\mathbf{x}]_{2r}$ already converge to the minimum $f_{\min}$ of f on $\mathbf{X}$ as $r \to \infty$ under mild conditions on $\mathbf{X}$ and μ [**Las11**]. (In the case that $\mathbf{X}$ is compact, it suffices that the reference measure μ is a finite Borel measure with support $\mathbf{X}$). The upper bounds (2.7) and (2.8) relying on the quadratic module and preordering of $\mathbf{X}$, which are more computationally intensive, are therefore not often studied in the literature. In this thesis, our focus will also be primarily on the bounds $\mathrm{ub}(f, \mathbf{X}, \mu)_r$, which rely simply on the sum-of-squares cone $\Sigma[\mathbf{x}]$.

The push-forward hierarchy. As we see below, the matrices involved in the computation of the bound $\mathrm{ub}(f, \mathbf{X}, \mu)_r$ are of size $\binom{n+r}{r}$. In an attempt to address this rapid growth, Lasserre [**Las20**] introduces a second type of upper bounds on $f_{\min}$ which are weaker but more economical. They provide a *univariate* approach to the problem by making use of *push-forward measures.* For a measure $\mu \in \mathcal{M}(\mathbf{X})$, the *push-forward* $\mu_f \in \mathcal{M}(\mathbb{R})$ of μ by f is defined by:

$$\mu_f(B) = \mu(f^{-1}(B)) \quad (B \subseteq \mathbb{R} \text{ Borel}).$$

Note that for any measurable function $g : \mathbb{R} \to \mathbb{R}$, we have:

$$\int_{f(\mathbf{X})} g(x) d\mu_f(x) = \int_{\mathbf{X}} g(f(\mathbf{x})) d\mu(\mathbf{x}).$$

For $r \in \mathbb{N}$, Lasserre [**Las20**] then defines an upper bound $\mathrm{ub}(f, \mathbf{X}, \mu)_r^{\mathrm{pf}}$ on $f_{\min}$ by:

$$\begin{aligned}\mathrm{ub}(f, \mathbf{X}, \mu)_r^{\mathrm{pf}} &:= \inf_{u \in \Sigma[x]_{2r}} \left\{ \int_{f(\mathbf{X})} x u(x) d\mu_f(x) : \int_{f(\mathbf{X})} u(x) d\mu_f(x) = 1 \right\} \\ &= \inf_{u \in \Sigma[x]_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) u(f(\mathbf{x})) d\mu(\mathbf{x}) : \int_{\mathbf{X}} u(f(\mathbf{x})) d\mu(\mathbf{x}) = 1 \right\}.\end{aligned}$$

The difference with the parameter $\mathrm{ub}(f, \mathbf{X}, \mu)_r$ is that we now restrict our search to *univariate* sums of squares $u \in \Sigma[x]_{2r}$. After composing such a polynomial u with f, we obtain a (multivariate) sum of squares $q = u \circ f \in \Sigma[\mathbf{x}]_{2rd}$ which is a feasible solution to (2.6) (of degree $d \cdot 2r$). Therefore, we have the inequality:

$$f_{\min} \leq \mathrm{ub}(f, \mathbf{X}, \mu)_{rd} \leq \mathrm{ub}(f, \mathbf{X}, \mu)_r^{\mathrm{pf}}.$$

Again, the parameter $\mathrm{ub}(f, \mathbf{X}, \mu)_r^{\mathrm{pf}}$ may be computed by solving a semidefinite program. The matrices involved, however, are now of much smaller size $r + 1$.

2.2. Semidefinite programming and sums of squares

There is a one-to-one correspondence between sums of squares of polynomials and positive semidefinite matrices. As we explain in this section, this correspondence allows us to compute the lower and upper bounds defined above using *semidefinite programming.* Our discussion here will be (very) brief, but there are many excellent and comprehensive resources on this topic, see, e.g., [**AL12, dKL19**].

Positive semidefinite matrices. A real, symmetric matrix $M \in \mathcal{S}^N$ of size $N \times N$ is called positive semidefinite (psd) if one of the following equivalent conditions hold:

$$\mathbf{x}^\top M \mathbf{x} \geq 0 \text{ for all } \mathbf{x} \in \mathbb{R}^N; \tag{2.9}$$

$$\text{The eigenvalues of } M \text{ are all greater or equal to } 0. \tag{2.10}$$

We also write $M \succeq 0$ when M is psd. If the conditions (2.9) or (2.10) hold with *strict* inequality, we say $M \succ 0$ is *strictly* psd. The set $\mathcal{S}_+^N$ of $N \times N$ positive semidefinite matrices is a *convex cone*, meaning that $\alpha A + \beta B \in \mathcal{S}_+^N$ for any two $A, B \in \mathcal{S}_+^N$ and scalars $\alpha, \beta \geq 0$. We have an inner product $\langle A, B \rangle := \mathrm{Tr}(AB)$ on $\mathcal{S}^N$, which allows us to define the so-called *dual cone* $(\mathcal{S}_+^N)^* := \{X \in \mathcal{S}^N : \langle X, M \rangle \geq 0 \text{ for all } M \in \mathcal{S}_+^N\}$. In fact, the cone $\mathcal{S}_+^N$ is *self-dual*, meaning that $\mathcal{S}_+^N = (\mathcal{S}_+^N)^*$.

Semidefinite programming. Let $C, A_i \in \mathcal{S}^N$ and $b_i \in \mathbb{R}$ for $i = 1, 2, \ldots, m$. An optimization problem of the form:

$$\text{val} = \sup_{X \in \mathcal{S}^N} \left\{ \langle C, X \rangle : \langle A_i, X \rangle = b_i \ (i = 1, 2, \ldots, m), \ X \succeq 0 \right\} \tag{P}$$

is called a *semidefinite program* (SDP). In principle, the program (P) could be *infeasible* (in which case val $:= -\infty$) or *unbounded* (in which case val $:= \infty$).

To each program (P), we associate the *dual program*:

$$\text{val}^* = \inf_{y \in \mathbb{R}^m} \left\{ \sum_{i=1}^m y_i b_i : \sum_{i=1}^m y_i A_i - C \succeq 0 \right\}. \tag{D}$$

The programs (P) and (D) satisfy *weak duality*, meaning that val $\leq$ val*. The reader familiar with linear programming might expect that they also satisfy *strong duality*, i.e., that val $=$ val*. In the case of semidefinite programming, however, this is not always the case. Fortunately, there are several sufficient conditions on (P) which guarantee that val $=$ val*. For instance, strong duality holds if (P) has a feasible solution $X \succ 0$ which is strictly psd (Slater's condition).

2.2.1. Formulation of the lower bounds as an SDP. We show how to reformulate the lower bounds defined above as semidefinite programs. As we will not make use of these reformulations directly in the rest of this thesis, we shall only give a rough sketch. The following key proposition links sums of squares of polynomials to positive semidefinite matrices.

PROPOSITION 2.4. *Let $q \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree $2r$. For $d \in \mathbb{N}$, we denote $\mathbb{N}^n_d := \{\alpha \in \mathbb{N}^n : \sum_{i=1}^n \alpha_i \leq d\}$. If we write $q(\mathbf{x}) = \sum_{\alpha \in \mathbb{N}^n_{2r}} q_\alpha \mathbf{x}^\alpha$ in the monomial basis, then q is a sum of squares if and only if there exists a psd matrix $Q = (Q_{\alpha,\beta})_{|\alpha|,|\beta| \leq r} \succeq 0$ such that:*

$$q(\mathbf{x}) = [\mathbf{x}]_r^\top Q [\mathbf{x}]_r = \langle [\mathbf{x}]_r [\mathbf{x}]_r^\top, Q \rangle = \sum_{\alpha, \beta \in \mathbb{N}^n_r} Q_{\alpha,\beta} \mathbf{x}^{\alpha+\beta}. \tag{2.11}$$

Here $[\mathbf{x}]_r = (\mathbf{x}^\alpha)_{|\alpha| \leq r}$ is the vector of monomials of degree at most r.

Thus one can check whether q is a sum of squares by checking whether a certain SDP is feasible. In light of Proposition 2.4, a polynomial $f(\mathbf{x}) = \sum_\gamma f_\gamma \mathbf{x}^\gamma$ has a representation in $\mathcal{Q}(\mathbf{X})_{2r}$ if and only if:

$$f(\mathbf{x}) = \sum_{j=0}^m g_j(\mathbf{x}) \langle [\mathbf{x}]_{r - \lceil \deg(g_j)/2 \rceil} [\mathbf{x}]_{r - \lceil \deg(g_j)/2 \rceil}^\top, Q_j \rangle$$

for appropriately sized matrices $Q_0, Q_1, \ldots, Q_m \succeq 0$. Therefore, after carefully selecting matrices $A_\gamma^{(j)}$ (depending on the constraints g_j defining $\mathbf{X}$), we may reformulate (2.3) as the SDP:

$$\sup_{Q_0,Q_1,\dots,Q_m\succeq 0} \Big\{ f_0 - \sum_{j=0}^{m} \langle A_0^{(j)}, Q_j\rangle : \sum_{j=0}^{m} \langle A_\gamma^{(j)}, Q_j\rangle = f_\gamma \quad \forall\, 1 \leq |\gamma| \leq 2r \Big\}.$$

The Schmüdgen-type bound (2.4) may be formulated as an SDP in a similar way, but there we would need a matrix Q_J for each $J \subseteq [m]$.

We note that in the literature, one often encounters the *dual* formulation of these SDPs instead (which are also known as *moment relaxations*). See for instance [**dKL19**],[**Las09b**]. In general, there can be a gap between the primal (sum-of-squares) and dual (moment) formulations. For semialgebraic sets $\mathbf{X}$ whose description satisfies an Archimedean condition, however, it can be shown that strong duality holds, see [**JH16**].

2.2.2. Formulation of the upper bounds as an SDP. Now, we show how to formulate the measure-based upper bounds as semidefinite programs. As the specifics of these formulations will play a role in the future, we will go into more detail here.

Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact set equipped with a finite, positive Borel measure μ supported on $\mathbf{X}$. For $\alpha \in \mathbb{N}^n$, the *moment* $m_\alpha \in \mathbb{R}$ of degree α of μ is defined as:

$$m_\alpha(\mu) := \int_{\mathbf{X}} \mathbf{x}^\alpha d\mu(\mathbf{x}).$$

For a polynomial $p \in \mathbb{R}[\mathbf{x}]$ and $r \in \mathbb{N}$, the (truncated) *moment matrix* $M_{p,r}(\mu)$ of p is then given by:

$$\big(M_{p,r}(\mu)\big)_{\alpha,\beta} := \int_{\mathbf{X}} p(\mathbf{x})\mathbf{x}^{\alpha+\beta} d\mu(\mathbf{x}) = \sum_{\gamma\in\mathbb{N}^n_r} p_\gamma m_{\alpha+\beta+\gamma}(\mu) \quad (\alpha,\beta \in \mathbb{N}^n_r). \tag{2.12}$$

In the case $p(\mathbf{x}) = 1$, we also write $M_r(\mu) = M_{1,r}(\mu)$ for simplicity. Now let $q \in \Sigma[\mathbf{x}]_{2r}$ be a sum of squares of degree r, and let $Q \succeq 0$ be the matrix corresponding to q in (2.11). Then we find that:

$$\int_{\mathbf{X}} p(\mathbf{x})q(\mathbf{x})d\mu(\mathbf{x}) = \langle M_{p,r}(\mu), Q\rangle. \tag{2.13}$$

Using relation (2.13) and Proposition 2.4, we may thus reformulate the upper bound $\mathrm{ub}(f,\mathbf{X},\mu)_r$ as the semidefinite program:

$$\mathrm{ub}(f,\mathbf{X},\mu)_r = \inf_{Q\succeq 0} \big\{ \langle M_{f,r}(\mu), Q\rangle : \langle M_r(\mu), Q\rangle = 1 \big\}. \tag{2.14}$$

Here, we optimize over matrices $Q \succeq 0$ of size $|\mathbb{N}^n_r| = \binom{n+r}{r}$.

Similarly, the push-forward bound $\mathrm{ub}(f,\mathbf{X},\mu)^{\mathrm{pf}}_r$ is given by:

$$\mathrm{ub}(f,\mathbf{X},\mu)^{\mathrm{pf}}_r = \inf_{Q\succeq 0} \big\{ \langle M_{x,r}(\mu_f), Q\rangle : \langle M_r(\mu_f), Q\rangle = 1 \big\}. \tag{2.15}$$

Since μ_f is a univariate measure, the optimization is now over matrices $Q \succeq 0$ of much smaller size $r+1$.

Computing moments. It is important to note that the reformulations (2.14) and (2.15) require knowledge of the moments $m_\alpha(\mu)$ and $m_k(\mu_f)$ of degree up to $2r + \deg(f)$, respectively. This is not a trivial requirement, as computing even just the *volume* of a semialgebraic set is hard in general [**Las09a**]. For many of the sets $\mathbf{X}$ and measures μ we consider, however, there are simple analytic expressions available for these moments. We discuss this further in Chapter 5.

An eigenvalue reformulation. The semidefinite programs (2.14) and (2.15) are special in the sense that they have only a single equality constraint. As we show now, this means that their dual formulations are in fact *generalized eigenvalue* problems. We only consider here the formulation (2.14) for the regular upper bounds, but the argument for (2.15) is the same. The dual program of (2.14) is given by:

$$\sup_{y \in \mathbb{R}} \big\{ y : M_{f,r}(\mu) - y M_r(\mu) \succeq 0 \big\}.$$

Using condition (2.10) for psd matrices, we see that the optimum value of this program is given by the smallest generalized eigenvalue of the system $\big(M_{f,r}(\mu), M_r(\mu)\big)$.

When $\mathbf{X} \subseteq \mathbb{R}^n$ is compact with non-empty interior, we have strong duality [**Las11**], meaning that:

$$\mathrm{ub}(f, \mathbf{X}, \mu)_r = \lambda_{\min}\big(M_{f,r}(\mu), M_r(\mu)\big). \tag{2.16}$$

Solving a generalized eigenvalue problem is much easier than solving an SDP (of the same size). The reformulation (2.16) thus reveals a potential computational advantage of the hierarchies of upper bounds over the hierarchies of lower bounds.

Orthonormal bases. In the above, we have always represented the space $\mathbb{R}[\mathbf{x}]_r$ of polynomials of degree at most r using the monomial basis. In principle, however, we could have used *any* basis. Indeed, Proposition 2.4 is true regardless of our choice of basis. Compared to (2.12), with respect to a general basis $\{P_\alpha : \alpha \in \mathbb{N}^n_r\}$, we get the moment matrices:

$$\big(M_{p,r}(\mu)\big)_{\alpha,\beta} := \int_{\mathbf{X}} p(\mathbf{x}) P_\alpha(\mathbf{x}) P_\beta(\mathbf{x}) d\mu(\mathbf{x}) \quad (\alpha, \beta \in \mathbb{N}^n_r). \tag{2.17}$$

The reformulation (2.14) is then exactly the same (but using these new moments matrices). One could also view this as a change of basis in the space of symmetric matrices in (2.14).

The point is that for a clever choice of $\{P_\alpha : \alpha \in \mathbb{N}^n_r\}$, the resulting program may be greatly simplified. In particular, if we choose the P_α to be an *orthonormal basis* w.r.t. $(\mathbf{X}, \mu)$, we find that $M_r(\mu) = I$ is the identity matrix, in which case (2.16) reduces to:

$$\text{ub}(f, \mathbf{X}, \mu)_r = \lambda_{\min}\big(M_{f,r}(\mu)\big). \tag{2.18}$$

As we see below in Section 2.3, this observation will sometimes permit an analysis of the upper bounds $\text{ub}(f, \mathbf{X}, \mu)_r$ when the eigenvalues of the matrix $M_{f,r}(\mu)$ are known.

2.3. Convergence analysis of the hierarchies

In this section, we discuss known results on the convergence rates of the hierarchies of lower and upper bounds defined above. An overview of all known (and new) results is given in Table 2.1 and Table 2.2 below.

2.3.1. The lower bounds. The first quantitative versions of Putinar's and Schmüdgen's Positivstellensätze are due to Nie and Schweighofer. For general Archimedean semialgebraic sets $\mathbf{X}$, they show that the Putinar-type bounds $\text{lb}(f, \mathcal{Q}(\mathbf{X}))_r$ converge to $f_{\min}$ at a rate in $O(1/\log(r)^c)$, where $c > 0$ is a constant depending on $\mathbf{X}$ [**NS07**]. For compact semialgebraic sets $\mathbf{X}$, Schweighofer [**Sch04**] shows that the Schmüdgen-type bounds $\text{lb}(f, \mathcal{T}(\mathbf{X}))_r$ converge to $f_{\min}$ at a rate in $O(1/r^c)$, where $c > 0$ is again a constant depending on $\mathbf{X}$. For a long time, these were the only general results available. In the very recent work [**BM21**], however, the authors show a convergence rate in $O(1/r^c)$ for the *Putinar*-type bounds on general Archimedean semialgebraic sets. They thus match the best known (general) rate for the Schmüdgen-type bounds and improve exponentially on the previous best known rate for $\text{lb}(f, \mathcal{Q}(\mathbf{X}))_r$ of [**NS07**].

Very roughly, the proofs of these general results rely on a clever embedding of the set $\mathbf{X}$ into a larger, simpler semialgebraic set, such as a box $[-1, 1]^n$ or a simplex Δ^n. The convergence rate of the bounds on $\mathbf{X}$ may then be analyzed in terms of the behaviour of the hierarchy on this simpler set. This is one motivation for studying the hierarchies for optimization over certain special, structured sets $\mathbf{X}$. Indeed our analysis [**LS21**] for the unit box $[-1, 1]^n$ (see below) is a key ingredient of the proof in [**BM21**].

Another motivation for studying special cases is that one may show much stronger guarantees there. For instance, in the case that $\mathbf{X} = [-1, 1]^n$ is the unit box, de Klerk and Laurent [**dKL10**] show a convergence rate in $O(1/r)$ for the Schmüdgen-type bounds. The same rate is shown by de Klerk and Kirschner [**KdK21**] when $\mathbf{X} = \Delta^n$ is the standard simplex. When $\mathbf{X} = S^{n-1}$ is the hypersphere, Fang and Fawzi [**FF21**] show a convergence rate in $O(1/r^2)$ for the Putinar-type bounds, which improves upon an earlier result in $O(1/r)$ due to Doherty and Wehner [**DW13**]. For the binary hypercube $\mathbf{X} = \{0, 1\}^n$ one may even show that the lower bound $\text{lb}(f, \mathcal{Q}(\mathbf{X}))_r$ is *exact* when $r \geq (n + d - 1)/2$ [**FSP16, STKI17**].

The proof techniques for the results in these special cases vary. For our purposes, the method Fang and Fawzi use for the analysis on S^{n-1} is most important, as it forms the basis of our results on the lower bounds in Part 2. We will discuss it in great detail in Chapter 6.

2.3.2. The upper bounds. Asymptotic convergence of the parameters $\mathrm{ub}(f, \mathbf{X}, \mu)_r$ and $\mathrm{ub}(f, \mathbf{X}, \mu)_r^{\mathrm{pf}}$ to $f_{\min}$ is shown by Lasserre in [**Las11**] and [**Las20**], respectively, under mild assumptions on $\mathbf{X}$ and μ. For the push-forward bounds, no quantitative results were known in the literature before our results in [**SL21a**], see below. On the other hand, rates for the parameter $\mathrm{ub}(f, \mathbf{X}, \mu)_r$ have been shown for different sets of assumptions on $\mathbf{X}, \mu$ and f. Depending on these assumptions, several strategies have been employed to obtain these rates, which we now discuss.

Algebraic analysis via an eigenvalue reformulation. The first strategy relies on the reformulation (2.18) of the optimization problem (2.6) as an eigenvalue minimization problem; particularly in the univariate case $n = 1$. Let $\{P_k \in \mathbb{R}[x]_k : k \in \mathbb{N}\}$ be the (unique) orthonormal basis of $\mathbb{R}[x]$ w.r.t. the inner product $\langle P_i, P_j \rangle = \int_{\mathbf{X}} P_i(x) P_j(x) d\mu(x)$. For $r \in \mathbb{N}$, we have seen that $\mathrm{ub}(f, \mathbf{X}, \mu)_r = \lambda_{\min}(M_{f,r}(\mu))$, which is the smallest eigenvalue of the matrix moment matrix $M_{f,r}(\mu)$ of (2.17). Any bounds on the eigenvalues of $M_{f,r}(\mu)$ thus immediately translate to bounds on $\mathrm{ub}(f, \mathbf{X}, \mu)_r$.

In [**dKL20b**], the authors determine the exact asymptotic behaviour of $\lambda_{\min}(M_{f,r}(\mu))$ in the case that f is a quadratic polynomial, $\mathbf{X} = [-1, 1]$ and $d\mu(x) = (1 - x^2)^{-\frac{1}{2}} dx$ is the Chebyshev measure. Based on this, they show that $\mathrm{ub}(f, \mathbf{X}, \mu)_r = O(1/r^2)$ and extend this result to arbitrary multivariate polynomials f on the hypercube $[-1, 1]^n$ equipped with the product measure $d\mu(x) = \prod_i^n (1 - x_i)^{-1/2} dx_i$. In addition, they prove that $\mathrm{ub}(f, \mathbf{X}, \mu)_r = \Theta(1/r^2)$ for linear polynomials, which thus shows that in some sense quadratic convergence is the best we can hope for.

For this latter result they make use of the fact that the moment matrix $M_{x,r}(\mu)$ for the linear polynomial $f(x) = x$ is precisely the matrix J defined in (1.3) for the orthogonal polynomials $\{P_k\}$. By Proposition 1.2, its smallest eigenvalue is thus given by the smallest root of the polynomial P_{r+1}, which in this case is the Chebyshev polynomial $\mathcal{C}_{r+1}$. The smallest root of $\mathcal{C}_{r+1}$ is known to converge to -1 at a rate in $\Theta(1/r^2)$, see Proposition 1.3.

The main disadvantage of the eigenvalue strategy is that it requires the moment matrix of f to have a closed form expression which is sufficiently structured so as to allow for an analysis of its eigenvalues. Closed form expressions for the entries of the matrix $M_{f,r}(\mu)$ are known only for special sets $\mathbf{X}$, such as the interval $[-1, 1]$, the unit ball, the unit sphere, or the simplex, and only with respect to certain measures.

However, as we will see in Chapter 3, the convergence analysis from [**dKL20b**] in $O(1/r^2)$ for the interval $[-1, 1]$ equipped with the Chebyshev measure can be transported to a large class of compact sets, such as the interval $[-1, 1]$ with more general measures, the ball, the simplex, and 'ball-like' convex bodies.

Analysis via the construction of feasible solutions. A second strategy to bound the convergence rate of the parameters $\mathrm{ub}(f, \mathbf{X}, \mu)_r$ is to construct explicit sum-of-squares density functions $q_r \in \Sigma[\mathbf{x}]_r$ for which the integral $\int_{\mathbf{X}} f(\mathbf{x}) q_r(\mathbf{x}) d\mu(\mathbf{x})$ is close to $f_{\min}$. In contrast to the previous strategy, such constructions will only yield upper bounds on the convergence rate of $\mathrm{ub}(f, \mathbf{X}, \mu)_r$.

As noted earlier, the integral $\int_{\mathbf{X}} f d\nu$ may be minimized by selecting the probability measure $\nu = \delta_{\mathbf{x}^*}$, i.e., the Dirac measure centered at a global minimizer $\mathbf{x}^*$ of f on $\mathbf{X}$. When the reference measure μ is the Lebesgue measure, it thus intuitively seems sensible to consider sum-of-squares densities q_r that approximate the Dirac delta in some way.

This approach is followed in [**dKLS17**]. There, the authors consider truncated Taylor expansions of the Gaussian function $e^{-x^2/2\sigma}$, which they use to define the sum-of-squares polynomials:

$$\phi_r(x) = \sum_{k=0}^{2r} \frac{1}{k!} \left(\frac{-x^2}{2\sigma} \right)^k \in \Sigma[x]_{2r} \qquad \text{for } r \in \mathbb{N}.$$

Setting $q_r(\mathbf{x}) \sim \phi_r(\|\mathbf{x} - \mathbf{x}^*\|)$ for carefully selected standard deviation $\sigma = \sigma(r)$, they show that $\int_{\mathbf{X}} f(\mathbf{x}) q_r(\mathbf{x}) d\mathbf{x} - f(\mathbf{x}^*) = O(1/\sqrt{r})$ when $\mathbf{X}$ satisfies a minor geometrical assumption (see Chapter 3), which holds, e.g., if $\mathbf{X}$ is a convex body or if it is star-shaped with respect to a ball.

In the subsequent work [**dKL18**], the authors show that if $\mathbf{X}$ is assumed to be a convex body, then a bound in $O(1/r)$ may be obtained by setting $q_r \sim \phi_r(f(\mathbf{x}))$. As explained in [**dKL18**], the sum-of-squares density q_r in this case can be seen as an approximation of the Boltzman density function for f, which plays an important role in simulated annealing.

The advantage of this second strategy appears to be its applicability to a broad class of sets $\mathbf{X}$ with respect to the natural Lebesgue measure. This generality, however, is offset by significantly weaker guarantees on $\mathrm{ub}(f, \mathbf{X}, \mu)_r$.

Analysis for the hypersphere. Tight results are known for polynomial minimization on the unit sphere $S^{n-1} = \{\mathbf{x} \in \mathbb{R}^n : \sum_i \mathbf{x}_i^2 = 1\}$, equipped with the uniform surface measure. Doherty and Wehner [**DW13**] have shown a convergence rate in $O(1/r)$, by using harmonic analysis on the sphere and connections to quantum information theory. In the recent work [**dKL20a**], the authors show an improved convergence rate in $O(1/r^2)$, by using a reduction to the case of the interval $[-1, 1]$ and the above mentioned convergence rate in $O(1/r^2)$ for this case. Such a reduction will form the basis of our arguments in Chapter 3. The reduction in [**dKL20a**] is based on replacing f by an easy (linear) upper estimator. This idea of using an upper estimator was already exploited in [**dKLS17, dKL20b**] (where a quadratic upper estimator was used) and we will also exploit it in Chapter 3.

Analysis using the Jackson kernel. Finally, we wish to mention a result of Hess, de Klerk and Laurent [**dKHL17**] for $\mathbf{X} = [-1,1]^n$. There, the authors show a convergence rate in $O(1/r^2)$ for the Schmüdgen-type upper bounds $\mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r$, when μ is the Chebyshev measure. Although their result is superceded by the later work [**dKL20b**] (which establishes a convergence rate in $O(1/r^2)$ for the weaker bounds $\mathrm{ub}(f, \mathbf{X}, \mu)_r$), their proof technique is of interest to us. Indeed, it relies on the well-known *Jackson kernel* (see, e.g., [**WWAF06**]), which will feature prominently in Chapter 9. We will also rely there on some technical properties of this kernel established in [**dKHL17**].

2.3.3. Negative results. To put the positive results on the convergence rates of the sum-of-squares hierarchies into perspective, let us summarize some *negative* results. First, as we mention above, one may show already for a linear polynomial f in the univariate setting that the upper bounds $\mathrm{ub}(f, [-1,1], \mu)_r$ converge at a rate in $\Omega(1/r^2)$ for a class of references measures μ (which includes the Lebesgue measure). This negative result carries over to a few other settings, such as the hypersphere. It relies either on the connection between the upper bounds and roots of orthogonal polynomials (see Section 2.2) or an a connection to cubature rules (see Section 6.3). In our estimation, it shows rather convincingly that a rate in $O(1/r^2)$ for the upper bounds is the best one should hope for in a general setting.

For the lower bounds, there is a quite extensive literature studying negative results on the binary cube $\{0,1\}^n$. We discuss some of these results in Chapter 7. There, the hierarchy is always exact at finite level $r \leq n$, and so one considers instead the dependence on n. In the setting where the hierarchy does not converge in finitely many steps, we are aware of only one example where the rate of convergence may be bounded from below. Namely, Stengle [**Ste96**] considers the function $f(x) = 1 - x^2$ on the interval $[-1,1]$ and shows that:

$$f_{\min} - \mathrm{lb}(f, \mathcal{Q}((1-x^2)^3)_r = \Omega(1/r^2).$$

It should be noted that his result relies on a nonstandard semialgebraic representation $[-1,1] = \{x \in \mathbb{R} : (1-x^2)^3 \geq 0\}$ of the interval (for the regular representation we would just have $1 - x^2 \in \mathcal{Q}(1-x^2)_2$). Nonetheless, it is quite interesting that he obtains a bound in $\Omega(1/r^2)$, which matches the best known rates in $O(1/r^2)$ on the hypersphere, the unit ball, the standard simplex and the unit cube $[-1,1]^n$ (see below).

2.4. Summary of results

To finish this chapter, we summarize the main new results on the hierarchies of upper and lower bounds proven in this thesis. See Table 2.1 and Table 2.2 for an overview.

2.4.1. The upper bounds. Our contributions for the upper bounds are presented in Part 1. First, in Chapter 3, we extend the best-possible convergence rate in $O(1/r^2)$ for the bounds $\mathrm{ub}(f, [-1, 1]^n)_r$ on the box $[-1, 1]$ w.r.t. the Chebyshev measure μ of [**dKL20b**] to a broader class of sets $\mathbf{X}$ and reference measures. This class includes the standard simplex, the unit ball, and 'ball-like' convex bodies w.r.t the Lebesgue measure. The main idea of our proof is to transport the result of [**dKL20b**] to these sets by showing that the behaviour of the upper bounds depends in essence only on the 'local behaviour' of the set $\mathbf{X}$, the measure μ and the polynomial f in the neighbourhood of a minimizer $\mathbf{x}^*$ of f.

Second, in Chapter 4, we establish a convergence rate in $O(\log^2 r/r^2)$ for the upper bounds $\mathrm{ub}(f, \mathbf{X})_r$ w.r.t. to the Lebesgue measure on a very large class of sets $\mathbf{X}$, which includes in particular all semialgebraic sets with dense interior. Our work shows a stronger rate and applies more generally than the earlier works [**dKLS17, dKL18**]. Indeed, the rate we show is only a log-factor away from best-possible. Somewhat surpisingly, our result in fact applies to the push-forward bounds $\mathrm{ub}(f, \mathbf{X})_r^{\mathrm{pf}}$, thereby giving the first analysis for this hierarchy. Our proof makes use of so-called *needle polynomials* to construct explicit sum-of-squares densities which approximate a Dirac function at a minimizer $\mathbf{x}^*$ of f on $\mathbf{X}$. The approximations we construct are better than those of [**dKLS17, dKL18**], thus yielding a stronger bound on the convergence rate.

2.4.2. The lower bounds. We present our new results on the lower bounds in Part 2. These results all rely on the same general proof technique, which is the subject of Chapter 6. Roughly speaking, this technique may be seen as a generalization of the one employed on the hypersphere in [**FF21**]. It relies on the *Christoffel-Darboux kernel* (aka *reproducing kernel*) and *Fourier analysis* to establish a link between the behaviour of the *lower* bounds on $\mathbf{X}$ and certain *univariate* instances of the *upper* bounds.

When $\mathbf{X}$ is sufficiently structured, this connection may then be exploited to obtain strong guarantees on the convergence rate of the lower bounds by analyzing (the roots of) classical *orthogonal polynomials*. We cover the binary cube $\{0, 1\}^n$ in Chapter 7. There, we make use of the *Funk-Hecke formula* to express the behaviour of the bounds in terms of roots of Krawtchouk polynomials, yielding an analysis in the regime $n \to \infty$ and $r = \Omega(n)$. The unit ball and standard simplex are treated in Chapter 8, where we rely on closed form expressions of the Christoffel-Darboux kernel in terms of Gegenbauer polynomials. Lastly, the case $\mathbf{X} = [-1, 1]^n$ is covered in Chapter 9. It relies on the aforementioned Jackson kernel and Chebyshev polynomials.

As we explain in Chapter 6, our analysis of the lower bounds in these settings also yields an analysis of the corresponding upper bounds. The obtained rates are included in the respective chapters and in Table 2.2 below. For the most part, however, they are superceded by earlier results, or by the results of Part 1.

Acknowledgments. We wish to thank Markus Schweighofer for bringing to our attention the negative result [**Ste96**] for the lower bounds on the interval.

X (compact)	**error**	**certificate**	**reference**
Archimedean	$O(1/\log(r)^c)$	$\mathcal{Q}(\mathbf{X})$	[**NS07**]
Archimedean	$O(1/r^c)$	$\mathcal{Q}(\mathbf{X})$	[**BM21**]
General	$O(1/r^c)$	$\mathcal{T}(\mathbf{X})$	[**Sch04**]
S^{n-1}	$O(1/r^2)$	$\mathcal{Q}(\mathbf{X})$ $(= \mathcal{T}(\mathbf{X}))$	[**FF21**]
$\{0,1\}^n$	see Theorem 7.1	$\mathcal{Q}(\mathbf{X})$ $(= \mathcal{T}(\mathbf{X}))$	Chapter 7
B^n	$O(1/r^2)$	$\mathcal{Q}(\mathbf{X})$ $(= \mathcal{T}(\mathbf{X}))$	Chapter 8
$[-1,1]^n$	$O(1/r^2)$	$\mathcal{T}(\mathbf{X})$	Chapter 9
Δ^n	$O(1/r)$	$\mathcal{T}(\mathbf{X})$	[**KdK21**]
Δ^n	$O(1/r^2)$	$\mathcal{T}(\mathbf{X})$	Chapter 8

TABLE 2.1. Overview of known and new results on the asymptotic error of Lasserre's hierarchies of lower bounds.

X (compact)	**error**	**certificate**	**measure**	**reference**
Geom. assumption	$O(1/\sqrt{r})$	$\Sigma[\mathbf{x}]$	Lebesgue	[**dKLS17**]
Convex body	$O(1/r)$	$\Sigma[\mathbf{x}]$	Lebesgue	[**dKL18**]
Semialg. + dense interior	$O(\log^2(r)/r^2)$	$\Sigma[\mathbf{x}]$	Lebesgue	Chapter 4
Semialg. + dense interior	$O(\log^2(r)/r^2)$	push-forward	Lebesgue	Chapter 4
S^{n-1}	$O(1/r^2)$	$\Sigma[\mathbf{x}]$	uniform	[**dKL20a**]
$[-1,1]^n$	$O(1/r^2)$	$\Sigma[\mathbf{x}]$	$\prod_i (1-\mathbf{x}_i)^\lambda d\mathbf{x}$ $(\lambda = -\frac{1}{2})$	[**dKL20b**]
$[-1,1]^n$	$O(1/r^2)$	$\Sigma[\mathbf{x}]$	$\prod_i (1-\mathbf{x}_i)^\lambda d\mathbf{x}$ $(\lambda \geq -\frac{1}{2})$	Chapter 3
$\{0,1\}^n$	see Theorem 7.3	$\Sigma[\mathbf{x}]$	uniform	Chapter 7
B^n	$O(1/r^2)$	$\Sigma[\mathbf{x}]$	$(1-\|\mathbf{x}\|^2)^\lambda d\mathbf{x}$ $(\lambda \geq 0)$	Chapter 3
Δ^n	$O(1/r^2)$	$\Sigma[\mathbf{x}]$	Lebesgue	Chapter 3
$[-1,1]^n$	$O(1/r^2)$	$\mathcal{T}(\mathbf{X})$	$\prod_i (1-x_i)^{-\frac{1}{2}} d\mathbf{x}$	[**dKHL17**]
B^n	$O(1/r^2)$	$\mathcal{T}(\mathbf{X})$	μ_B, see (8.16)	Chapter 8
Δ^n	$O(1/r^2)$	$\mathcal{T}(\mathbf{X})$	μ_Δ, see (8.18)	Chapter 8

TABLE 2.2. Overview of known and new results on the asymptotic error of Lasserre's hierarchies of upper bounds.

Part 1

Measure-based hierarchies of upper bounds

CHAPTER 3

Convergence analysis for measure-based bounds I

The worst is yet to come.

Arthur Schopenhauer

This chapter is based on my joint work [**SL20**] *with Monique Laurent.*

Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact set equipped with a finite, positive Borel measure μ supported on $\mathbf{X}$, and let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree d. Recall the measure-based hierarchy of upper bounds on the minimum $f_{\min}$ of f on $\mathbf{X}$ introduced in Chapter 2:

$$\mathrm{ub}(f, \mathbf{X}, \mu)_r := \inf_{q \in \Sigma[\mathbf{x}]_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mu(\mathbf{x}) : \int_{\mathbf{X}} q(\mathbf{x}) d\mu(\mathbf{x}) = 1 \right\}. \tag{3.1}$$

De Klerk and Laurent [**dKL20b**] show a convergence rate of the bounds (3.1) in $O(1/r^2)$ when $\mathbf{X} = [-1, 1]^n$ is the unit box, equipped with the Chebyshev measure $d\mu(\mathbf{x}) = \prod_{i=1}^n (1 - \mathbf{x}_i^2)^{-\frac{1}{2}} d\mathbf{x}_i$. That is, they show that:

$$\mathrm{Error}(f; \mathbf{X}, \mu)_r := \mathrm{ub}(f, \mathbf{X}, \mu)_r - f_{\min} = O(1/r^2)$$

for this choice of $\mathbf{X}$ and reference measure μ. In fact, they show that this rate is best-possible already when f has degree $d = 1$.

Outline. In this chapter, we extend their result to a broader class of convex bodies $\mathbf{X} \subseteq \mathbb{R}^n$ and reference measures μ. First, in Section 3.3, we show that for the hypercube $\mathbf{X} = [-1, 1]^n$, we have convergence in $O(1/r^2)$ for $f \in \mathbb{R}[\mathbf{x}]$ of arbitrary degree and all measures of the form $d\mu(\mathbf{x}) = \prod_{i=1}^n (1 - \mathbf{x}_i^2)^{\lambda} d\mathbf{x}_i$ with $\lambda > -1/2$. Of particular interest is the case $\lambda = 0$, where we have the Lebesgue measure on $[-1, 1]^n$. Next, in Section 3.4, we use this result to show convergence in $O(1/r^2)$ of the measure-based bounds on the unit ball B^n for all references measures of the form $d\mu(\mathbf{x}) = (1 - \|\mathbf{x}\|^2)^{\lambda} dx$ with $\lambda \geq 0$. We then apply this result in Section 3.5 to prove a rate in $O(1/r^2)$ when μ is the Lebesgue measure and $\mathbf{X}$ is a 'ball-like' convex body, meaning roughly that it has inscribed and circumscribed tangent balls at all boundary points (see Definition 3.22 below). Finally, we consider the standard simplex $\Delta^n \subseteq \mathbb{R}^n$ in Section 3.6, where we also obtain a rate in $O(1/r^2)$ for the Lebesgue measure.

The primary new tool we use to obtain these results is Proposition 3.12 (see Section 3.2), which tells us that the asymptotic behaviour of the measure-based bounds essentially only depends on the *local* geometry of $\mathbf{X}$ in a neighbourhood of a global minimizer $\mathbf{x}^* \in \mathbf{X}$ of f, and the behaviour of f and μ in this neighbourhood. This tool allows us to transport the result of de Klerk and Laurent on $[-1,1]^n$ to the new sets.

3.1. Preliminaries

We first introduce some notation that we will use throughout the rest of this chapter and recall some basic terminology and results about convex bodies. Then, we cover some basic techniques that will simplify the analysis of the measure-based bounds in later sections. Finally, we discuss some properties of measures and weight functions needed to prove Proposition 3.12 below.

3.1.1. Notation. For $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, $\langle \mathbf{x}, \mathbf{y} \rangle$ denotes the standard inner product of $\mathbf{x}$ and $\mathbf{y}$, and $\|\mathbf{x}\|^2 = \langle \mathbf{x}, \mathbf{x} \rangle$ the corresponding norm. We write $B^n_\rho(c) := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x} - c\| \leq \rho\}$ for the n-dimensional ball of radius ρ centered at $c \in \mathbb{R}^n$. When $\rho = 1$ and $c = 0$, we also use our usual notation $B^n = B^n_1(0)$.

Throughout, $\mathbf{X} \subseteq \mathbb{R}^n$ is always a compact set with non-empty interior, and f is an n-variate polynomial. We let $\nabla f(\mathbf{x})$ (resp., $\nabla^2 f(\mathbf{x})$) denote the gradient (resp., the Hessian) of f at $\mathbf{x} \in \mathbb{R}^n$, and introduce the parameters:

$$\Gamma_{\text{grad}}(f, \mathbf{X}) := \max_{\mathbf{x} \in \mathbf{X}} \|\nabla f(\mathbf{x})\| \quad \text{and} \quad \Gamma_{\text{hess}}(f, \mathbf{X}) := \frac{1}{2} \max_{\mathbf{x} \in \mathbf{X}} \|\nabla^2 f(\mathbf{x})\|. \tag{3.2}$$

Here, $\|\cdot\|$ denotes the Euclidean norm. Whenever we write an expression of the form:

$$\text{“Error}(f; \mathbf{X}, \mu)_r = O(1/r^2)\text{”},$$

we mean that there exists a constant $c > 0$ such that $\text{Error}(f; \mathbf{X}, \mu)_r \leq c/r^2$ for all $r \in \mathbb{N}$, where c depends only on $\mathbf{X}, \mu$, and the parameters $\Gamma_{\text{grad}}(f, \mathbf{X}), \Gamma_{\text{hess}}(f, \mathbf{X})$. Some of our results are obtained by embedding $\mathbf{X}$ into a larger set $\widehat{\mathbf{X}} \subseteq \mathbb{R}^n$. If this is the case, then c may depend on $\Gamma_{\text{grad}}(f, \widehat{\mathbf{X}}), \Gamma_{\text{hess}}(f, \widehat{\mathbf{X}})$ as well. If there is an additional dependence of c on the global minimizer $\mathbf{x}^*$ of f on $\mathbf{X}$, we will make this explicit by using the notation “$O_{\mathbf{x}^*}$”.

3.1.2. Convex bodies. Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a convex body, i.e., a compact, convex set with non-empty interior. We say $v \in \mathbb{R}^n$ is an (inward) normal of $\mathbf{X}$ at $a \in \mathbf{X}$ if $\langle v, \mathbf{x} - a \rangle \geq 0$ holds for all $\mathbf{x} \in \mathbf{X}$. We refer to the set of all normals of $\mathbf{X}$ at a as the normal cone , and write

$$N_{\mathbf{X}}(a) := \{v \in \mathbb{R}^n : \langle v, \mathbf{x} - a \rangle \geq 0 \text{ for all } \mathbf{x} \in \mathbf{X}\}.$$

We will make use of the following basic result.

LEMMA 3.1 (e.g., [**BL06**, Prop. 2.1.1]). *Let $\mathbf{X}$ be a convex body and let $g : \mathbb{R}^n \to \mathbb{R}$ be a continuously differentiable function with local minimizer $\mathbf{x}^* \in \mathbf{X}$. Then $\nabla g(\mathbf{x}^*) \in N_{\mathbf{X}}(\mathbf{x}^*)$.*

PROOF. Suppose not. Then, by definition of $N_{\mathbf{X}}(\mathbf{x}^*)$, there exists an element $\mathbf{y} \in \mathbf{X}$ such that $\langle \nabla g(\mathbf{x}^*), \mathbf{y} - \mathbf{x}^* \rangle < 0$. Expanding the definition of the gradient this means that

$$0 > \langle \nabla g(\mathbf{x}^*), \mathbf{y} - \mathbf{x}^* \rangle = \lim_{t \downarrow 0} \frac{g(t\mathbf{y} + (1-t)\mathbf{x}^*) - g(\mathbf{x}^*)}{t},$$

which implies $g(t\mathbf{y} + (1-t)\mathbf{x}^*) < g(\mathbf{x}^*)$ for all $t > 0$ small enough. But $t\mathbf{y} + (1-t)\mathbf{x}^* \in \mathbf{X}$ by convexity, contradicting the fact that $\mathbf{x}^*$ is a local minimizer of g on $\mathbf{X}$. □

The set $\mathbf{X}$ is *smooth* if it has a unique unit normal $v(a)$ at each boundary point $a \in \partial\mathbf{X}$. In this case, we denote by $T_a\mathbf{X}$ the (unique) hyperplane tangent to $\mathbf{X}$ at a, defined by the equation $\langle \mathbf{x} - a, v(a) \rangle = 0$.

For $k \geq 1$, we say $\mathbf{X}$ is of class C^k if there exists a convex function $\Psi \in C^k(\mathbb{R}^n, \mathbb{R})$ such that $\mathbf{X} = \{\mathbf{x} \in \mathbb{R}^n : \Psi(x) \leq 0\}$ and $\partial\mathbf{X} = \{\mathbf{x} \in \mathbb{R}^n : \Psi(\mathbf{x}) = 0\}$. If $\mathbf{X}$ is of class C^k for some $k \geq 1$, it is automatically smooth in the above sense.

We refer, e.g., to [**BF87**] for a general reference on convex bodies.

3.1.3. Linear transformations. Suppose that $\phi : \mathbb{R}^n \to \mathbb{R}^n$ is a nonsingular affine transformation, given by $\phi(\mathbf{x}) = U\mathbf{x} + a$. If q is a sum-of-squares density function w.r.t. the Lebesgue measure on $\phi(\mathbf{X})$, then we have:

$$\int_{\phi(\mathbf{X})} q(y) f(\phi^{-1}(\mathbf{y})) d\mathbf{y} = |\det U| \cdot \int_{\mathbf{X}} q(\phi(\mathbf{x})) f(\mathbf{x}) d\mathbf{x} \quad \text{and}$$

$$1 = \int_{\phi(\mathbf{X})} q(\mathbf{y}) d\mathbf{y} = |\det U| \cdot \int_{\mathbf{X}} q(\phi(\mathbf{x})) d\mathbf{x}.$$

As a result, the polynomial $\widehat{q} := (q \circ \phi) / \int_{\mathbf{X}} q(\phi(\mathbf{x})) d\mathbf{x} = (q \circ \phi) \cdot |\det U|$ is a sum of squares density function w.r.t. the Lebesgue measure on $\mathbf{X}$. It has the same degree as q, and it satisfies:

$$\int_{\mathbf{X}} \widehat{q}(\mathbf{x}) f(\mathbf{x}) d\mathbf{x} = \int_{\phi(\mathbf{X})} q(\mathbf{x}) f(\phi^{-1}(\mathbf{x})) d\mathbf{x}.$$

We have just shown the following.

LEMMA 3.2. *Let $\phi : \mathbb{R}^n \to \mathbb{R}^n$ be a non-singular affine transformation. Write $g := f \circ \phi^{-1}$. Then we have*

$$\mathrm{Error}(f; \mathbf{X})_r = \mathrm{Error}(g; \phi(\mathbf{X}))_r.$$

3.1.4. Upper estimators. Given a point $a \in \mathbf{X}$ and two functions $f, g : \mathbf{X} \to \mathbb{R}$, we write $f \leq_a g$ if $f(a) = g(a)$ and $f(\mathbf{x}) \leq g(\mathbf{x})$ for all $\mathbf{x} \in \mathbf{X}$; we then say that g is an *upper estimator* for f on $\mathbf{X}$, which is *exact at* a. The next lemma, whose easy proof is omitted, will be very useful.

LEMMA 3.3. *Let $g : \mathbf{X} \to \mathbb{R}$ be an upper estimator for f, exact at one of the global minimizers of f on $\mathbf{X}$. Then we have* $\mathrm{Error}(f; \mathbf{X}, \mu)_r \leq \mathrm{Error}(g; \mathbf{X}, \mu)_r$ *for all $r \in \mathbb{N}$.*

REMARK 3.4. *We make the following observations for future reference.*

1. *Lemma 3.3 tells us that we may always replace f in our analysis by an upper estimator which is exact at one of its global minimizers. This is useful if we can find an upper estimator that is significantly simpler to analyze.*
2. *We may always assume that $f_{\min} = 0$, in which case $f(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in \mathbf{X}$ and $\mathrm{Error}(f; \mathbf{X}, \mu)_r = \mathrm{ub}(f, \mathbf{X}, \mu)_r$. Indeed, if we consider the function g given by $g(\mathbf{x}) = f(\mathbf{x}) - f_{\min}$, then $g_{\min,\mathbf{X}} = 0$, and for every density function q on $\mathbf{X}$, we have:*
$$\int_{\mathbf{X}} g(\mathbf{x})q(\mathbf{x})d\mu(\mathbf{x}) = \int_{\mathbf{X}} f(\mathbf{x})q(\mathbf{x})d\mu(\mathbf{x}) - f_{\min},$$
showing that $\mathrm{Error}(f; \mathbf{X}, \mu)_r = \mathrm{ub}(g, \mathbf{X}, \mu)_r$ for all $r \in \mathbb{N}$.

In the remainder of this section, we derive some general upper estimators based on the following variant of Taylor's theorem for multivariate functions.

THEOREM 3.5 (Taylor's theorem). *For $f \in C^2(\mathbb{R}^n, \mathbb{R})$ and $a \in \mathbf{X}$ we have:*
$$f(\mathbf{x}) \leq f(a) + \langle \nabla f(a), \mathbf{x} - a \rangle + \Gamma_{\text{hess}}(\mathbf{X}, f)\|\mathbf{x} - a\|^2 \quad \text{for all } \mathbf{x} \in \mathbf{X},$$
where $\Gamma_{\text{hess}}(\mathbf{X}, f)$ is the constant from (3.2).

LEMMA 3.6. *Let $\mathbf{x}^* \in \mathbf{X}$ be a global minimizer of f on $\mathbf{X}$. Then f has an upper estimator g on $\mathbf{X}$ which is exact at $\mathbf{x}^*$ and satisfies the following properties:*

(i) *g is a quadratic, separable polynomial.*
(ii) *$g(\mathbf{x}) \geq f(\mathbf{x}^*) + \Gamma_{\text{hess}}(\mathbf{X}, f)\|\mathbf{x} - \mathbf{x}^*\|^2$ for all $\mathbf{x} \in \mathbf{X}$.*
(iii) *If $\mathbf{x}^* \in \mathrm{int}(\mathbf{X})$, then $g(\mathbf{x}) \leq f(\mathbf{x}^*) + \Gamma_{\text{hess}}(\mathbf{X}, f)\|\mathbf{x} - \mathbf{x}^*\|^2$ for all $\mathbf{x} \in \mathbf{X}$.*

PROOF. Consider the function g defined by:
$$g(\mathbf{x}) := f(\mathbf{x}^*) + \langle \nabla f(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle + \Gamma_{\text{hess}}(\mathbf{X}, f)\|\mathbf{x} - \mathbf{x}^*\|^2, \tag{3.3}$$
which is an upper estimator of f exact at $\mathbf{x}^*$ by Theorem 3.5. As we have $\|\mathbf{x} - \mathbf{x}^*\|^2 = \sum_i^n (\mathbf{x}_i - \mathbf{x}_i^*)^2$, g is indeed a quadratic, separable polynomial.

As $\mathbf{x}^*$ is a global minimizer of f on $\mathbf{X}$, we know by Lemma 3.1 that $\nabla f(\mathbf{x}^*) \in N_{\mathbf{X}}(\mathbf{x}^*)$. This means that $\langle \nabla f(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle \geq 0$ for all $\mathbf{x} \in \mathbf{X}$, which proves the second property.

If $\mathbf{x}^* \in \mathrm{int}(\mathbf{X})$, we must have $\nabla f(\mathbf{x}^*) = 0$, and the third property follows. □

In the special case that $\mathbf{X}$ is a ball and f has a global minimizer $\mathbf{x}^*$ on the boundary of $\mathbf{X}$, we have an upper estimator for f, exact at $\mathbf{x}^*$, which is a linear polynomial.

LEMMA 3.7. *Assume that $f(\mathbf{x}^*) = f_{\min, B_\rho^n(c)}$ for some $\mathbf{x}^* \in \partial B_\rho^n(c)$. Then there exists a linear polynomial g with $f \leq_{\mathbf{x}^*} g$ on $B_\rho^n(c)$.*

PROOF. Write $\mathbf{X} = B^n_\rho(c)$ and $\gamma = \Gamma_{\text{grad}}(f, \mathbf{X})$ for simplicity. In view of Lemma 3.6, we have $f(\mathbf{x}) \leq g(\mathbf{x})$ for all $\mathbf{x} \in \mathbf{X}$, where g is the quadratic polynomial from relation (3.3). Since $\mathbf{x}^* \in \partial\mathbf{X}$ is a global minimizer of f on $\mathbf{X}$, we have $\nabla f(\mathbf{x}^*) \in N_{\mathbf{X}}(\mathbf{x}^*)$ by Lemma 3.1, and thus $\nabla f(\mathbf{x}^*) = \lambda(c - \mathbf{x}^*)$ for some $\lambda \geq 0$. Therefore we have:

$$\langle \nabla f(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle = \langle \lambda(c - \mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle = \lambda\rho^2 + \lambda\langle \mathbf{x} - c, c - \mathbf{x}^* \rangle.$$

On the other hand, for any $\mathbf{x} \in \mathbf{X}$ we have:

$$\|\mathbf{x} - \mathbf{x}^*\|^2 = \|\mathbf{x} - c\|^2 + \|c - \mathbf{x}^*\|^2 + 2\langle \mathbf{x} - c, c - \mathbf{x}^* \rangle \leq 2\rho^2 + 2\langle \mathbf{x} - c, c - \mathbf{x}^* \rangle.$$

Combining these facts we get:

$$f(\mathbf{x}) \leq g(\mathbf{x}) \leq f(\mathbf{x}^*) + (\lambda + 2\gamma)(\rho^2 + \langle \mathbf{x} - c, c - \mathbf{x}^* \rangle) =: h(\mathbf{x}).$$

So $h(\mathbf{x})$ is a linear upper estimator of f with $h(\mathbf{x}^*) = f(\mathbf{x}^*)$, as desired.

□

REMARK 3.8. *As can be seen from the above proof, the assumption in Lemma 3.7 that $\mathbf{x}^* \in \partial\mathbf{X} = \partial B^n_\rho(c)$ is a global minimizer of f on $\mathbf{X}$ may be replaced by the weaker assumption that $\nabla f(\mathbf{x}^*) \in N_{\mathbf{X}}(\mathbf{x}^*)$.*

3.1.5. Measures and weight functions. A function $w : \text{int}(\mathbf{X}) \to \mathbb{R}_{>0}$ is a *weight function* on $\mathbf{X}$ if it is continuous and satisfies $0 < \int_{\mathbf{X}} w(\mathbf{x})d\mathbf{x} < \infty$. A weight function w gives rise to a measure μ_w on $\mathbf{X}$ defined by $d\mu_w(\mathbf{x}) := w(\mathbf{x})d\mathbf{x}$. We note that if $\mathbf{X} \subseteq \widehat{\mathbf{X}}$, and $\widehat{w}$ is a weight function on $\widehat{\mathbf{X}}$, it can naturally be interpreted as a weight function on $\mathbf{X}$ as well, by simply restricting its domain (assuming $\int_{\mathbf{X}} \widehat{w}(\mathbf{x})d\mathbf{x} > 0$). In what follows we will implicitly make use of this fact.

DEFINITION 3.9. *Given two weight functions $w, \widehat{w}$ on $\mathbf{X}$ and a point $a \in \mathbf{X}$, we say that $\widehat{w} \preceq_a w$ on $\mathbf{X}$ if there exist constants $\varepsilon, m_a > 0$ such that*

$$m_a\widehat{w}(\mathbf{x}) \leq w(\mathbf{x}) \text{ for all } \mathbf{x} \in B^n_\varepsilon(a) \cap \text{int}(\mathbf{X}). \tag{3.4}$$

If the constant m_a can be chosen uniformly, i.e., if there exists a constant $m > 0$ such that

$$m\widehat{w}(\mathbf{x}) \leq w(\mathbf{x}) \text{ for all } \mathbf{x} \in \text{int}(\mathbf{X}), \tag{3.5}$$

then we say that $\widehat{w} \preceq w$ on $\mathbf{X}$.

REMARK 3.10. *We note the following facts for future reference:*

- (i) *As weight functions are continuous on the interior of $\mathbf{X}$ by definition, we always have $\widehat{w} \preceq_a w$ if $a \in \text{int}(\mathbf{X})$.*
- (ii) *If w is bounded from below, and $\widehat{w}$ is bounded from above on $\text{int}(\mathbf{X})$, then we automatically have $\widehat{w} \preceq w$.*

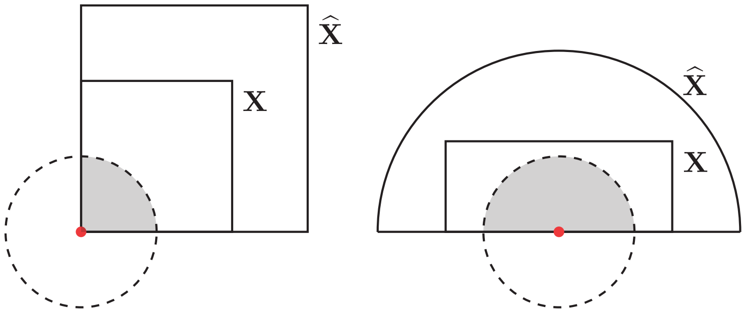

FIGURE 3.1. Some examples of sets $\mathbf{X}, \widehat{\mathbf{X}}$ for which $\mathbf{X} \subseteq_a \widehat{\mathbf{X}}$. The red dot indicates the point a, and the gray area indicates $B^n_\varepsilon(a) \cap \mathbf{X}$.

3.2. Local similarity

Assuming that the global minimizer $\mathbf{x}^*$ of f on $\mathbf{X}$ is unique, sum-of-squares density functions q for which the integral $\int_{\mathbf{X}} q(\mathbf{x})f(\mathbf{x})d\mu(\mathbf{x})$ is small should in some sense approximate the Dirac delta function centered at $\mathbf{x}^*$. With this in mind, it seems reasonable to expect that the quality of the bound $\text{ub}(f, \mathbf{X}, \mu)_r$ depends in essence only on the *local* properties of f, $\mathbf{X}$ and μ around $\mathbf{x}^*$. In this section, we formalize this intuition.

DEFINITION 3.11. *Suppose* $\mathbf{X} \subseteq \widehat{\mathbf{X}} \subseteq \mathbb{R}^n$. *Given* $a \in \mathbf{X}$, *we say that* $\mathbf{X}$ *and* $\widehat{\mathbf{X}}$ *are* locally similar *at* a, *which we denote by* $\mathbf{X} \subseteq_a \widehat{\mathbf{X}}$, *if there exists* $\varepsilon > 0$ *such that*

$$B^n_\varepsilon(a) \cap \mathbf{X} = B^n_\varepsilon(a) \cap \widehat{\mathbf{X}}.$$

Clearly, $\mathbf{X} \subseteq_a \widehat{\mathbf{X}}$ *for any point* $a \in \text{int}(\mathbf{X})$.

Figure 3.1 depicts some examples of locally similar sets.

PROPOSITION 3.12. *Let* $\mathbf{X} \subseteq \widehat{\mathbf{X}} \subseteq \mathbb{R}^n$, *let* $\mathbf{x}^* \in \mathbf{X}$ *be a global minimizer of* f *on* $\mathbf{X}$ *and assume* $\mathbf{X} \subseteq_{\mathbf{x}^*} \widehat{\mathbf{X}}$. *Let* $w, \widehat{w}$ *be two weight functions on* $\mathbf{X}, \widehat{\mathbf{X}}$, *respectively. Assume that* $\widehat{w}(\mathbf{x}) \geq w(\mathbf{x})$ *for all* $\mathbf{x} \in \text{int}(\mathbf{X})$, *and that* $\widehat{w} \preceq_{\mathbf{x}^*} w$. *Then there exists an upper estimator* g *of* f *on* $\widehat{\mathbf{X}}$ *which is exact at* $\mathbf{x}^*$ *and satisfies*

$$\text{Error}(g; \mathbf{X}, w)_r \leq \frac{2}{m_{\mathbf{x}^*}} \text{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r$$

for all $r \in \mathbb{N}$ *large enough. Here* $m_{\mathbf{x}^*} > 0$ *is the constant defined by* (3.4).

Recall that if g is an upper estimator for f which is exact at one of its global minimizers, we then have $\text{Error}(f; \mathbf{X}, w)_r \leq \text{Error}(g; \mathbf{X}, w)_r$ by Lemma 3.3. Proposition 3.12 then allows us to bound $\text{Error}(f; \mathbf{X}, w)_r$ in terms of $\text{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r$. For its proof, we first need the following lemma.

LEMMA 3.13. *Let $a \in \mathbf{X}$, and assume that $\mathbf{X} \subseteq_a \widehat{\mathbf{X}}$. Then any normal vector of $\mathbf{X}$ at a is also a normal vector of $\widehat{\mathbf{X}}$. That is, $N_{\mathbf{X}}(a) \subseteq N_{\widehat{\mathbf{X}}}(a)$.*

PROOF. Let $v \in N_{\mathbf{X}}(a)$. Suppose for contradiction that $v \notin N_{\widehat{\mathbf{X}}}(a)$. Then, by definition of the normal cone, there exists $\mathbf{y} \in \widehat{\mathbf{X}}$ such that $\langle v, \mathbf{y} - a\rangle < 0$. As $\mathbf{X} \subseteq_a \widehat{\mathbf{X}}$, there exists $\varepsilon > 0$ for which $\mathbf{X} \cap B_a^n(\varepsilon) = \widehat{\mathbf{X}} \cap B_a^n(\varepsilon)$. Now choose $1 > \eta > 0$ small enough such that $\mathbf{y}' := \eta\mathbf{y} + (1-\eta)a \in B_a^n(\varepsilon)$. Then, by convexity, we have $\mathbf{y}' \in \widehat{\mathbf{X}} \cap B_a^n(\varepsilon) = \mathbf{X} \cap B_a^n(\varepsilon)$. Now, we have $\langle v, \mathbf{y}' - a\rangle = \eta\langle v, \mathbf{y} - a\rangle < 0$. But, as $\mathbf{y}' \in \mathbf{X}$, this contradicts the assumption that $v \in N_{\mathbf{X}}(a)$. □

PROOF OF PROPOSITION 3.12. For simplicity, we assume here $f(\mathbf{x}^*) = 0$, which is without loss of generality by Remark 3.4. Consider the quadratic polynomial g from (3.3):

$$g(\mathbf{x}) = \langle \nabla f(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^*\rangle + \gamma\|\mathbf{x} - \mathbf{x}^*\|^2,$$

where $\gamma := \Gamma_{\mathrm{hess}}(\widehat{\mathbf{X}}, f)$ is defined in (3.2). By Taylor's theorem (Theorem 3.5), we have that $g(\mathbf{x}) \geq f(\mathbf{x})$ for all $\mathbf{x} \in \widehat{\mathbf{X}}$, and clearly $g(\mathbf{x}^*) = f(\mathbf{x}^*)$. That is, g is an upper estimator for f on $\widehat{\mathbf{X}}$, exact at $\mathbf{x}^*$ (cf. Lemma 3.6). We proceed to show that:

$$\mathrm{Error}(g; \mathbf{X}, w)_r \leq \frac{2}{m_{\mathbf{x}^*}} \mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r.$$

We start by selecting a degree $2r$ sum-of-squares polynomial $\widehat{q_r}$ satisfying

$$\int_{\widehat{\mathbf{X}}} \widehat{q_r}(\mathbf{x})\widehat{w}(\mathbf{x})d\mathbf{x} = 1 \quad \text{and} \quad \int_{\widehat{\mathbf{X}}} g(\mathbf{x})\widehat{q_r}(\mathbf{x})\widehat{w}(\mathbf{x})d\mathbf{x} = \mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r.$$

We may then rescale $\widehat{q_r}$ to obtain a density function $q_r \in \Sigma_{2r}$ on $\mathbf{X}$ w.r.t. w by setting:

$$q_r := \frac{\widehat{q_r}}{\int_{\mathbf{X}} \widehat{q_r}(\mathbf{x})w(\mathbf{x})d\mathbf{x}}.$$

By assumption, $w(\mathbf{x}) \leq \widehat{w}(\mathbf{x})$ for all $\mathbf{x} \in \mathrm{int}(\mathbf{X})$. Moreover, $g(\mathbf{x}) \geq f(\mathbf{x}^*) = 0$ for all $\mathbf{x} \in \mathrm{int}(\mathbf{X})$. This implies that:

$$\begin{aligned}\mathrm{Error}(g; \mathbf{X}, w)_r &\leq \int_{\mathbf{X}} g(\mathbf{x})q_r(\mathbf{x})w(\mathbf{x})d\mathbf{x} \\ &\leq \frac{\int_{\widehat{\mathbf{X}}} g(\mathbf{x})\widehat{q_r}(\mathbf{x})\widehat{w}(\mathbf{x})d\mathbf{x}}{\int_{\mathbf{X}} \widehat{q_r}(\mathbf{x})w(\mathbf{x})d\mathbf{x}} = \frac{\mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r}{\int_{\mathbf{X}} \widehat{q_r}(\mathbf{x})w(\mathbf{x})d\mathbf{x}}\end{aligned}$$

and thus it suffices to show that $\int_{\mathbf{X}} \widehat{q_r}(\mathbf{x})w(\mathbf{x})d\mathbf{x} \geq \frac{1}{2}m_{\mathbf{x}^*}$. The key to proving this bound is the following lemma, which tells us that optimum sum-of-squares densities should assign rather high weight to the ball $B_\varepsilon^n(\mathbf{x}^*)$ around $\mathbf{x}^*$.

LEMMA 3.14. *Let $\varepsilon > 0$. Then, for any $r \in \mathbb{N}$, we have:*

$$\int_{B_\varepsilon^n(a)\cap\widehat{\mathbf{X}}} \widehat{q_r}(\mathbf{x})\widehat{w}(\mathbf{x})d\mathbf{x} \geq 1 - \frac{\mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r}{\gamma\varepsilon^2}.$$

Proof. By Lemma 3.1, we have $\nabla f(a) \in N_{\mathbf{X}}(a)$ and so $\nabla f(a) \in N_{\widehat{\mathbf{X}}}(a)$ by Lemma 3.13. As a result, we have $g(\mathbf{x}) \geq \gamma\|\mathbf{x}-a\|^2$ for all $\mathbf{x} \in \widehat{\mathbf{X}}$ (cf. Lemma 3.6). In particular, this implies that $g(\mathbf{x}) \geq \gamma\|\mathbf{x}-a\|^2 \geq \gamma\varepsilon^2$ for all $\mathbf{x} \in \widehat{\mathbf{X}} \setminus B^n_\varepsilon(a)$ and so:

$$\begin{aligned}\text{Error}(g;\widehat{\mathbf{X}},\widehat{w})_r &\geq \int_{\widehat{\mathbf{X}}\setminus B^n_\varepsilon(a)} g(\mathbf{x})\widehat{q_r}(\mathbf{x})\widehat{w}(\mathbf{x})d\mathbf{x}\\ &\geq \gamma\varepsilon^2 \int_{\widehat{\mathbf{X}}\setminus B^n_\varepsilon(a)} \widehat{q_r}(\mathbf{x})\widehat{w}(\mathbf{x})d\mathbf{x}\\ &= \gamma\varepsilon^2\left(1-\int_{B^n_\varepsilon(a)\cap\widehat{\mathbf{X}}} \widehat{q_r}(\mathbf{x})\widehat{w}(\mathbf{x})d\mathbf{x}\right).\end{aligned}$$

The statement now follows from reordering terms. □

As $\mathbf{X} \subseteq_a \widehat{\mathbf{X}}$, there exists $\varepsilon_1 > 0$ such that $B^n_{\varepsilon_1}(a) \cap \mathbf{X} = B^n_{\varepsilon_1}(a) \cap \widehat{\mathbf{X}}$. As $\widehat{w} \preceq_a w$, there exist $\varepsilon_2 > 0$, $m_a > 0$ such that $m_a\widehat{w}(\mathbf{x}) \leq w(\mathbf{x})$ for $\mathbf{x} \in B^n_{\varepsilon_2}(a) \cap \text{int}(\mathbf{X})$. Set $\varepsilon = \min\{\varepsilon_1, \varepsilon_2\}$. Choose $r_0 \in \mathbb{N}$ large enough such that $\text{Error}(g;\widehat{\mathbf{X}},\widehat{w})_r < \frac{\varepsilon^2\gamma}{2}$ for all $r \geq r_0$, which is possible since $\text{Error}(g;\widehat{\mathbf{X}},\widehat{w})_r$ tends to 0 as $r \to \infty$. Then, Lemma 3.14 yields:

$$\int_{B^n_\varepsilon(a)\cap\widehat{\mathbf{X}}} \widehat{q_r}(\mathbf{x})\widehat{w}(\mathbf{x})d\mathbf{x} \geq \frac{1}{2}$$

for all $r \geq r_0$. Putting things together, we obtain the desired lower bound:

$$\begin{aligned}\int_{\mathbf{X}} \widehat{q_r}(\mathbf{x})w(\mathbf{x})d\mathbf{x} &\geq \int_{B^n_\varepsilon(a)\cap\mathbf{X}} \widehat{q_r}(\mathbf{x})w(\mathbf{x})d\mathbf{x}\\ &\geq m_a \int_{B^n_\varepsilon(a)\cap\widehat{\mathbf{X}}} \widehat{q_r}(\mathbf{x})\widehat{w}(\mathbf{x})d\mathbf{x} \geq \frac{1}{2}m_a.\end{aligned}$$

for all $r \geq r_0$. □

Corollary 3.15. *Let $\mathbf{X} \subseteq \widehat{\mathbf{X}} \subseteq \mathbb{R}^n$, let $\mathbf{x}^* \in \mathbf{X}$ be a global minimizer of f on $\mathbf{X}$, and assume that $\mathbf{X} \subseteq_{\mathbf{x}^*} \widehat{\mathbf{X}}$. Let $w, \widehat{w}$ be two weight functions on $\mathbf{X}, \widehat{\mathbf{X}}$, respectively. Assume that $\widehat{w}(\mathbf{x}) \geq w(\mathbf{x})$ for all $\mathbf{x} \in \text{int}(\mathbf{X})$ and that $\widehat{w} \preceq w$. Then there exists an upper estimator g of f on $\widehat{\mathbf{X}}$, exact at $\mathbf{x}^*$, such that*

$$\text{Error}(g;\mathbf{X},w)_r \leq \frac{2}{m}\text{Error}(g;\widehat{\mathbf{X}},\widehat{w})_r$$

for all $r \in \mathbb{N}$ large enough. Here $m > 0$ is the constant defined by (3.5).

3.3. The unit cube

Here we consider optimization over the hypercube $\mathbf{X} = [-1,1]^n$ and we restrict to reference measures on $\mathbf{X}$ having a weight function of the form

$$\widehat{w_\lambda}(\mathbf{x}) := \prod_{i=1}^n w_\lambda(\mathbf{x}_i) = \prod_{i=1}^n (1-\mathbf{x}_i^2)^\lambda \tag{3.6}$$

with $\lambda > -1$. The following result is shown in [**dKL20b**] on the convergence rate of the bound $\text{Error}(f;\mathbf{X},\widehat{w_\lambda})_r$ when using the measure $\widehat{w_\lambda}(\mathbf{x})d\mathbf{x}$ on $\mathbf{X} = [-1,1]^n$.

THEOREM 3.16 ([**dKL20b**]). *Let $\mathbf{X} = [-1,1]^n$ and consider the weight function $\widehat{w_\lambda}$ from* (3.6).

(i) *If $\lambda = -\frac{1}{2}$, then we have:*

$$\text{Error}(f;\mathbf{X},\widehat{w_\lambda})_r = O\left(\frac{1}{r^2}\right). \tag{3.7}$$

(ii) *If $n = 1$ and f has a global minimizer on the boundary of* $[-1,1]$, *then* (3.7) *holds for all $\lambda > -1$.*

The key ingredients for claim (ii) above are: (a) when the global minimizer is a boundary point of $[-1,1]$ then f has a linear upper estimator (recall Lemma 3.7), and (b) the convergence rate of (3.7) holds for any linear function and any $\lambda > -1$ (see Section 2.3 and[**dKL20b**]).

In this section we show Theorem 3.17 below, which extends the above result to all weight functions $\widehat{w_\lambda}(\mathbf{x})$ with $\lambda \geq -\frac{1}{2}$. Following the approach in [**dKL20b**], we proceed in two steps: first we reduce to the univariate case, and then we deal with the univariate case. Then the new situation to be dealt with is when $n = 1$ and the minimizer lies in the interior of $[-1,1]$, which we can settle by getting back to the case $\lambda = -\frac{1}{2}$ through applying Proposition 3.12, the 'local similarity' tool, with $\mathbf{X} = \widehat{\mathbf{X}} = [-1,1]$.

3.3.1. Reduction to the univariate case. Let $\mathbf{x}^* \in \mathbf{X}$ be a global minimizer of f in $\mathbf{X} = [-1,1]^n$. Following [**dKL20b**] (recall Remark 3.4 and Lemma 3.6), we consider the upper estimator $f(\mathbf{x}) \leq_a g(\mathbf{x}) := f(a) + \langle \nabla f(a), \mathbf{x} - a\rangle + \gamma_{f,\mathbf{X}}\|\mathbf{x} - a\|^2$. This g is separable, i.e., we can write $g(\mathbf{x}) = \sum_{i=1}^n g_i(\mathbf{x}_i)$, where each g_i is quadratic univariate with a_i as global minimizer over $[-1,1]$. Let q_r^i be an optimum solution to the problem (3.1) corresponding to the minimization of g_i over $[-1,1]$ w.r.t. the weight function $w_\lambda(\mathbf{x}_i) = (1-\mathbf{x}_i^2)^\lambda$. If we set $q_r(\mathbf{x}) = \prod_{i=1}^n q_r^i(\mathbf{x}_i)$, then q_r is a sum of squares with degree at most $2nr$, such that $\int_{\mathbf{X}} q_r(\mathbf{x})\widehat{w_\lambda}(\mathbf{x})d\mathbf{x} = 1$. Hence we have:

$$\begin{aligned}
\mathrm{ub}(f, \mathbf{X}, \widehat{w_\lambda})_{rn} - f(\mathbf{x}^*) &\leq \int_{\mathbf{X}} f(\mathbf{x})q_r(\mathbf{x})\widehat{w_\lambda}(\mathbf{x})d\mathbf{x} - f(\mathbf{x}^*)\\
&\leq \int_{\mathbf{X}} g(\mathbf{x})q_r(\mathbf{x})\widehat{w_\lambda}(\mathbf{x})d\mathbf{x} - g(\mathbf{x}^*)\\
&= \sum_{i=1}^{n} \Big(\int_{-1}^{1} g_i(\mathbf{x})q_r^i(\mathbf{x}_i)w_\lambda(\mathbf{x}_i)d\mathbf{x}_i - g_i(\mathbf{x}_i^*) \Big)\\
&= \sum_{i=1}^{n} \big(\mathrm{ub}(g_i, [-1,1], w_\lambda)_r - g_i(\mathbf{x}_i^*)\big)\\
&= \sum_{i=1}^{n} \mathrm{Error}(g_i; [-1,1], w_\lambda)_r.
\end{aligned}$$

As a consequence, we need only to consider the case of a quadratic univariate polynomial f on $\mathbf{X} = [-1, 1]$. We distinguish two cases, depending whether the global minimizer lies on the boundary or in the interior of $\mathbf{X}$. The case where the global minimizer lies on the boundary of $[-1, 1]$ is settled by Theorem 3.16(ii) above, so we next assume the global minimizer lies in the interior of $[-1, 1]$.

Case of a global minimizer $\mathbf{x}^*$ in the interior of $\mathbf{X} = [-1, 1]$. To deal with this case we make use of Proposition 3.12 with $\mathbf{X} = \widehat{\mathbf{X}} = [-1, 1]$, weight function $w(\mathbf{x}) := w_\lambda(\mathbf{x})$ on $\mathbf{X}$, and weight function $\widehat{w}(\mathbf{x}) := w_{-1/2}(\mathbf{x})$ on $\widehat{\mathbf{X}}$. We check that the conditions of the proposition are met. As $\widehat{\mathbf{X}} = \mathbf{X}$, clearly we have $\mathbf{X} \subseteq_{\mathbf{x}^*} \widehat{\mathbf{X}}$. Further, for any $\lambda \geq -\frac{1}{2}$, we have

$$w_\lambda(\mathbf{x}) = (1 - \mathbf{x}^2)^\lambda \leq (1 - \mathbf{x}^2)^{-\frac{1}{2}} = w_{-1/2}(\mathbf{x})$$

for all $\mathbf{x} \in (-1, 1) = \mathrm{int}(\mathbf{X})$. As $\mathbf{x}^* \in \mathrm{int}(\mathbf{X})$, we also have $w_\lambda \preceq_{\mathbf{x}^*} w_{-1/2}$ (see Remark 3.10(i)). Hence we may apply Proposition 3.12 to find that there exists a polynomial upper estimator g of f on $[-1, 1]$, exact at $\mathbf{x}^*$, and having

$$\mathrm{Error}(g; \mathbf{X}, w)_r \leq \frac{2}{m_a}\mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r$$

for all $r \in \mathbb{N}$ large enough. Now, (the univariate case of) Theorem 3.16(i) allows us to claim $\mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r = O(1/r^2)$, so that we obtain:

$$\mathrm{Error}(f; \mathbf{X}, w_\lambda)_r \leq \mathrm{Error}(g; \mathbf{X}, w_\lambda)_r = O_{\mathbf{x}^*}\big(\mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r\big) = O_{\mathbf{x}^*}(1/r^2).$$

In summary, in view of the above, we have shown the following extension of Theorem 3.16.

THEOREM 3.17. *Let* $\mathbf{X} = [-1,1]^n$ *and* $\lambda \geq -\frac{1}{2}$. *Let* $\mathbf{x}^*$ *be a global minimizer of* f *on* $\mathbf{X}$. *Then we have*

$$\text{Error}(f; \mathbf{X}, \widehat{w_\lambda})_r = O_{\mathbf{x}^*}\left(\frac{1}{r^2}\right).$$

The constant $m_{\mathbf{x}^*}$ involved in the proof of Theorem 3.17 depends on the global minimizer $\mathbf{x}^*$ of f on $[-1,1]$. It is introduced by the application of Proposition 3.12 to cover the case where $\mathbf{x}^*$ lies in the interior of $[-1,1]$. When $\lambda = 0$ (i.e., when $w = w_0 = 1$ corresponds to the Lebesgue measure), one can replace $m_{\mathbf{x}^*}$ by a uniform constant $m > 0$, as we now explain.

Consider $\widehat{\mathbf{X}} := [-2,2] \supseteq [-1,1] = \mathbf{X}$, equipped with the scaled Chebyshev weight $\widehat{w}(\mathbf{x}) := w_{-1/2}(\mathbf{x}/2) = (1-\mathbf{x}^2/4)^{-1/2}$. Of course, Theorem 3.16 applies to this choice of $\widehat{\mathbf{X}}, \widehat{w}$ as well. Further, we still have $\widehat{w}(\mathbf{x}) \geq w(\mathbf{x}) = w_0(\mathbf{x}) = 1$ for all $\mathbf{x} \in [-1,1]$. However, we now have a *uniform* upper bound $\widehat{w}(\mathbf{x}) \leq \widehat{w}(1)$ for $\widehat{w}$ on $\mathbf{X}$, which means that $\widehat{w} \preceq w$ on $\mathbf{X}$ (see Remark 3.10(ii)). Indeed, we have:

$$\widehat{w}(\mathbf{x})/\widehat{w}(1) \leq 1 = w_0(\mathbf{x}) = w(\mathbf{x}) \quad \text{for all } \mathbf{x} \in [-1,1].$$

We may thus apply Corollary 3.15 (instead of Proposition 3.12) to obtain the following.

COROLLARY 3.18. *If* $\mathbf{X} = [-1,1]^n$ *is equipped with the Lebesgue measure, then:*

$$\text{Error}(f; \mathbf{X})_r = O\left(\frac{1}{r^2}\right).$$

3.4. The unit ball

We now consider optimization over the unit ball $\mathbf{X} = B^n \subseteq \mathbb{R}^n$ ($n \geq 2$); we restrict to reference measures on B^n with weight function of the form:

$$w_\lambda(\mathbf{x}) = (1 - \|\mathbf{x}\|^2)^\lambda, \tag{3.8}$$

where $\lambda > -1$. For further reference we recall (see e.g. [**DX14**, §6.3.2]) or [**DX13**, §11]) that

$$C_{n,\lambda} := \int_{B^n} w_\lambda(\mathbf{x}) d\mathbf{x} = \frac{\pi^{\frac{n}{2}} \Gamma(\lambda+1)}{\Gamma\left(\lambda + 1 + \frac{n}{2}\right)}. \tag{3.9}$$

For the case $\lambda \geq 0$, we can analyse the bounds and show the following result.

THEOREM 3.19. *Let* $\mathbf{X} = B^n$ *be the unit ball. Let* $\mathbf{x}^*$ *be a global minimizer of* f *on* $\mathbf{X}$. *Consider the weight function* w_λ *from* (3.8) *on* $\mathbf{X}$.

(i) *If* $\lambda = 0$, *we have*

$$\text{Error}(f; \mathbf{X}, w_\lambda)_r = O\left(\frac{1}{r^2}\right).$$

(ii) *If $\lambda > 0$, we have*

$$\mathrm{Error}(f; \mathbf{X}, w_\lambda)_r = O_{\mathbf{x}^*}\left(\frac{1}{r^2}\right).$$

For the proof, we distinguish the two cases where $\mathbf{x}^*$ lies in the interior of $\mathbf{X}$ or on its boundary.

3.4.1. Case of a global minimizer in the interior of X. Our strategy is to reduce this to the case of the hypercube with the help of Proposition 3.12. Set $\widehat{\mathbf{X}} := [-1,1]^n \supseteq B^n = \mathbf{X}$. As $\mathbf{x}^* \in \mathrm{int}(\mathbf{X})$, we have $\mathbf{X} \subseteq_{\mathbf{x}^*} \widehat{\mathbf{X}}$. Consider the weight function $w(\mathbf{x}) := w_\lambda(\mathbf{x}) = (1 - \|\mathbf{x}\|^2)^\lambda$ on $\mathbf{X}$, and $\widehat{w}(\mathbf{x}) := 1$ on the hypercube $\widehat{\mathbf{X}}$. Since $\lambda \geq 0$, we have $w_\lambda(\mathbf{x}) \leq 1 \leq \widehat{w}(\mathbf{x})$ for all $\mathbf{x} \in \mathbf{X}$. Furthermore, as $\mathbf{x}^* \in \mathrm{int}(\mathbf{X})$, we also have $\widehat{w} \preceq_a w$. Hence we may apply Proposition 3.12 to find a polynomial upper estimator g of f on $\widehat{\mathbf{X}}$, exact at $\mathbf{x}^*$, satisfying:

$$\mathrm{Error}(g; \mathbf{X}, w)_r \leq \frac{2}{m_{\mathbf{x}^*}} \mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r$$

for all $r \in \mathbb{N}$ large enough. Here $m_{\mathbf{x}^*} > 0$ is the constant from (3.4). Now, Theorem 3.17 allows us to claim $\mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r = O_{\mathbf{x}^*}(1/r^2)$. Hence we obtain:

$$\mathrm{Error}(f; \mathbf{X}, w)_r \leq \mathrm{Error}(g; \mathbf{X}, w)_r = O_{\mathbf{x}^*}(\mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r) = O_{\mathbf{x}^*}(1/r^2).$$

As in the previous section, it is possible to replace the constant $m_{\mathbf{x}^*}$ by a uniform constant $m > 0$ in the case that $\lambda = 0$, i.e., in the case that we have the Lebesgue measure on $\mathbf{X}$. Indeed, in this case we have $\widehat{w} = w$ $(= w_0 = 1)$, and so in particular $\widehat{w} \preceq w$. We may thus invoke Corollary 3.15 (instead of Proposition 3.12) to obtain

$$\mathrm{Error}(g; \mathbf{X}, w)_r \leq 2\mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r$$

and so

$$\mathrm{Error}(f; \mathbf{X}, w)_r = O(\mathrm{Error}(g; \widehat{\mathbf{X}}, \widehat{w})_r) = O(1/r^2).$$

Note that in this case, we do not actually make use of the fact that $\mathbf{X} = B^n$. Rather, we only need that $\mathbf{x}^*$ lies in the interior of $\mathbf{X}$ and that $\mathbf{X} \subseteq [-1,1]^n$. As we may freely apply affine transformations to $\mathbf{X}$ (by Lemma 3.2), the latter is no true restriction. We have thus shown the following result.

THEOREM 3.20. *Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact set, with non-empty interior, equipped with the Lebesgue measure. Assume that f has a global minimizer $\mathbf{x}^*$ on $\mathbf{X}$ with $\mathbf{x}^* \in \mathrm{int}(\mathbf{X})$. Then we have:*

$$\mathrm{Error}(f; \mathbf{X})_r = O\left(\frac{1}{r^2}\right).$$

3.4.2. Case of a global minimizer on the boundary of X. Our strategy is now to reduce to the univariate case of the interval $[-1,1]$. For this, we use Lemma 3.7, which claims that f has a linear upper estimator g on $\mathbf{X}$, exact at $\mathbf{x}^*$. Up to applying an orthogonal transformation (and scaling) we may assume that g is of the form $g(\mathbf{x}) = \mathbf{x}_1$. It therefore suffices now to analyze the behaviour of the bounds for the function $\mathbf{x}_1$ minimized on the ball B^n. Note that when minimizing $\mathbf{x}_1$ on B^n or on the interval $[-1,1]$ the minimum is attained at the boundary in both cases. The following technical lemma will be useful for reducing to the case of the interval $[-1,1]$.

LEMMA 3.21. *Let h be a univariate polynomial and let $\lambda > -1$. Then we have*

$$\int_{B^n} h(\mathbf{x}_1) w_\lambda(\mathbf{x}) d\mathbf{x} = C_{n-1,\lambda} \int_{-1}^{1} h(\mathbf{x}_1) w_{\lambda+\frac{n-1}{2}}(\mathbf{x}_1) d\mathbf{x}_1,$$

where $C_{n-1,\lambda}$ is given in (3.9).

PROOF. Change variables and set $u_j = \frac{\mathbf{x}_j}{\sqrt{1-\mathbf{x}_1^2}}$ for $2 \leq j \leq d$. Then we have

$$w_\lambda(\mathbf{x}) = (1 - \mathbf{x}_1^2 - \mathbf{x}_2^2 + \ldots - \mathbf{x}_n^2)^\lambda = (1-\mathbf{x}_1^2)^\lambda (1 - u_2^2 - \ldots - u_n^2)^\lambda$$

and $d\mathbf{x}_2 \cdots d\mathbf{x}_n = (1-\mathbf{x}_1^2)^{\frac{n-1}{2}} du_2 \cdots du_n$. Putting things together we obtain the desired result. $\square$

Let $q_r(\mathbf{x}_1)$ be an optimal sum-of-squares density with degree at most $2r$ for the problem of minimizing $\mathbf{x}_1$ over the interval $[-1,1]$, equipped with the weight function $w(\mathbf{x}) := w_{\lambda+\frac{n-1}{2}}(\mathbf{x})$. Then, its scaling $C_{n-1,\lambda}^{-1} q_r(\mathbf{x}_1)$ provides a feasible solution for the problem of minimizing $g(\mathbf{x}) = \mathbf{x}_1$ over the ball $\mathbf{X} = B^n$. Indeed, using Lemma 3.21, we have $\int_{B^n} C_{n-1,\lambda}^{-1} q_r(\mathbf{x}_1) w_\lambda(\mathbf{x}) d\mathbf{x} = \int_{-1}^{1} q_r(\mathbf{x}_1) w(\mathbf{x}) d\mathbf{x}_1 = 1$, and so

$$g_{\mathbf{X},w_\lambda}^{(r)} \leq \int_{B^n} \mathbf{x}_1 C_{n-1,\lambda}^{-1} q_r(\mathbf{x}_1) w_\lambda(\mathbf{x}) d\mathbf{x} = \int_{-1}^{1} \mathbf{x}_1 q_r(\mathbf{x}_1) w(\mathbf{x}_1) d\mathbf{x}_1.$$

The proof is now concluded by applying Theorem 3.16(ii).

3.5. Ball-like convex bodies

Here we show a convergence rate of $\mathrm{Error}(f;\mathbf{X})_r$ in $O(1/r^2)$ for a special class of smooth convex bodies $\mathbf{X}$ with respect to the Lebesgue measure. The basis for this result is a reduction to the case of the unit ball.

We say $\mathbf{X}$ has an *inscribed tangent ball* (of radius ε) at $\mathbf{x} \in \partial\mathbf{X}$ if there exists $\varepsilon > 0$ and a closed ball B_{insc} of radius ε such that $\mathbf{x} \in \partial B_{\text{insc}}$ and $B_{\text{insc}} \subseteq \mathbf{X}$. Similarly, we say $\mathbf{X}$ has a *circumscribed tangent ball* (of radius ε) at $\mathbf{x} \in \partial\mathbf{X}$ if there exists $\varepsilon > 0$ and a closed ball B_{circ} of radius ε such that $\mathbf{x} \in \partial B_{\text{circ}}$ and $\mathbf{X} \subseteq B_{\text{circ}}$.

DEFINITION 3.22. *We say that a (smooth) convex body* $\mathbf{X}$ *is* ball-like *if there exist (uniform)* $\varepsilon_{\text{insc}}, \varepsilon_{\text{circ}} > 0$ *such that* $\mathbf{X}$ *has inscribed and circumscribed tangent balls of radii* $\varepsilon_{\text{insc}}, \varepsilon_{\text{circ}}$*, respectively, at all points* $\mathbf{x} \in \partial\mathbf{X}$*.*

THEOREM 3.23. *Assume that* $\mathbf{X}$ *is a (smooth) ball-like convex body, equipped with the Lebesgue measure. Then we have*

$$\text{Error}(f; \mathbf{X})_r = O\left(\frac{1}{r^2}\right).$$

PROOF. Let $\mathbf{x}^* \in \mathbf{X}$ be a global minimizer of f on $\mathbf{X}$. We again distinguish two cases depending on whether $\mathbf{x}^*$ lies in the interior of $\mathbf{X}$ or on its boundary; the case of a global minimizer in the interior of $\mathbf{X}$ is covered directly by Theorem 3.20.

3.5.1. Case of a global minimizer on the boundary of X. By applying a suitable affine transformation, we can arrange that the following holds: $f(\mathbf{x}^*) = 0$, $\mathbf{x}^* = 0$, e_1 is an inward normal of $\mathbf{X}$ at a, and the radius of the circumscribed tangent ball B_{circ} at $\mathbf{x}^*$ is equal to 1, i.e., $B_{\text{circ}} = B_1^n(e_1)$. See Figure 3.2 for an illustration. Now, as a is a global minimizer of f on $\mathbf{X}$, we have $\nabla f(\mathbf{x}^*) \in N_{\mathbf{X}}(\mathbf{x}^*)$ by Lemma 3.1. But $N_{\mathbf{X}}(\mathbf{x}^*) = N_{B_{\text{circ}}}(\mathbf{x}^*)$, and so $\nabla f(\mathbf{x}^*) \in N_{B_{\text{circ}}}(\mathbf{x}^*)$. As noted in Remark 3.8, we may thus use Lemma 3.7 to find that $f(\mathbf{x}) \leq_{\mathbf{x}^*} c\langle e_1, \mathbf{x}\rangle = c\mathbf{x}_1$ on B_{circ} for some constant $c > 0$. In light of Remark 3.4(i), and after scaling, it therefore suffices to analyse the function $f(\mathbf{x}) = \mathbf{x}_1$.

Again, we will use a reduction to the univariate case, now on the interval $[0, 2]$. For any $r \in \mathbb{N}$, let $q_r \in \Sigma[x]_{2r}$ be an optimum sum-of-squares density of degree $2r$ for the minimization of $\mathbf{x}_1$ on $[0, 2]$ with respect to the weight function:

$$w'(\mathbf{x}_1) := w_{\frac{n-1}{2}}(\mathbf{x}_1 - 1) = [1 - (\mathbf{x}_1 - 1)^2]^{\frac{n-1}{2}} = [2\mathbf{x}_1 - \mathbf{x}_1^2]^{\frac{n-1}{2}}.$$

That is, $q_r \in \Sigma[x]_{2r}$ satisfies:

$$\int_0^2 \mathbf{x}_1 q_r(\mathbf{x}_1) w'(\mathbf{x}_1) d\mathbf{x}_1 = O(1/r^2) \quad \text{and} \quad \int_0^2 q_r(\mathbf{x}_1) w'(\mathbf{x}_1) d\mathbf{x}_1 = 1, \quad (3.10)$$

where the first equality relies on Theorem 3.16(ii). As

$$\mathbf{x} \mapsto q_r(\mathbf{x}_1)/(\int_{\mathbf{X}} q_r(\mathbf{x}_1) d\mathbf{x})$$

is a sum-of-squares density on $\mathbf{X}$ with respect to the Lebesgue measure, we have:

$$\text{Error}(f; \mathbf{X})_r \leq \frac{\int_{\mathbf{X}} \mathbf{x}_1 q_r(\mathbf{x}_1) d\mathbf{x}}{\int_{\mathbf{X}} q_r(\mathbf{x}_1) d\mathbf{x}}. \quad (3.11)$$

We will now show that, on the one hand, the numerator $\int_{\mathbf{X}} \mathbf{x}_1 q_r(\mathbf{x}_1) d\mathbf{x}$ in (3.11) has an upper bound in $O(1/r^2)$ and that, on the other hand, the denominator $\int_{\mathbf{X}} q_r(\mathbf{x}_1) d\mathbf{x}$ in (3.11) is lower bounded by an absolute constant that does not depend on r. Putting these two bounds together then yields $\mathrm{Error}(f;\mathbf{X})_r = O(1/r^2)$, as desired.

The upper bound. We make use of the fact that $\mathbf{X} \subseteq B_{\mathrm{circ}}$ to compute:

$$\begin{aligned}
\int_{\mathbf{X}} \mathbf{x}_1 q_r(\mathbf{x}_1) d\mathbf{x} &\leq \int_{B_{\mathrm{circ}}} \mathbf{x}_1 q_r(\mathbf{x}_1) d\mathbf{x} \\
&= \int_{B^n} (\mathbf{y}_1 + 1) q_r(\mathbf{y}_1 + 1) d\mathbf{y} && [\mathbf{y} = \mathbf{x} - e_1] \\
&= C_{n-1,0} \int_{-1}^{1} (\mathbf{y}_1 + 1) q_r(\mathbf{y}_1 + 1) w_{\frac{n-1}{2}}(\mathbf{y}_1) d\mathbf{y}_1 && [\text{by Lemma 3.21}] \\
&= C_{n-1,0} \int_{0}^{2} z q_r(z) w'(z) dz && [z = \mathbf{y}_1 + 1] \\
&= O(1/r^2). && [\text{by (3.10)}]
\end{aligned}$$

The lower bound. Here, we consider an inscribed tangent ball B_{insc} of $\mathbf{X}$ at $\mathbf{x}^* = 0$. Say $B_{\mathrm{insc}} = B^n_\rho(\rho e_1)$ for some $\rho > 0$. See again Figure 3.2. We may then compute:

$$\begin{aligned}
\int_{\mathbf{X}} q_r(\mathbf{x}_1) d\mathbf{x} &\geq \int_{B_{\mathrm{insc}}} q_r(\mathbf{x}_1) d\mathbf{x} \\
&= \int_{B^n} q_r\big(\rho(\mathbf{y}_1 + 1)\big) \rho^n d\mathbf{y} && [\mathbf{y} = \frac{\mathbf{x} - \rho e_1}{\rho}] \\
&= \rho^n C_{n-1,0} \int_{-1}^{1} q_r\big(\rho(\mathbf{y}_1 + 1)\big) w_{\frac{n-1}{2}}(\mathbf{y}_1) d\mathbf{y}_1 && [\text{by Lemma 3.21}] \\
&= \rho^{n-1} C_{n-1,0} \int_{0}^{2\rho} q_r(z) w_{\frac{n-1}{2}}(z/\rho - 1) dz && [z = \rho(\mathbf{y}_1 + 1)] \\
&\geq \rho^{n-1} C_{n-1,0} \int_{0}^{\rho} q_r(z) w'(z) \frac{w_{\frac{n-1}{2}}(z/\rho - 1)}{w_{\frac{n-1}{2}}(z - 1)} dz && [w'(z) = w_{\frac{n-1}{2}}(z-1)] \\
&\geq \left(\frac{\rho}{2-\rho}\right)^{\frac{n-1}{2}} C_{n-1,0} \int_{0}^{\rho} q_r(z) w'(z) dz,
\end{aligned}$$

where the last inequality follows using the fact that $\frac{1-(z/\rho-1)^2}{1-(z-1)^2} \geq \frac{1}{\rho(2-\rho)}$ for $z \in [0, \rho]$. It remains to show that

$$\int_0^\rho q_r(z) w'(z) dz \geq \frac{1}{2} \quad \text{for all } r \text{ large enough.}$$

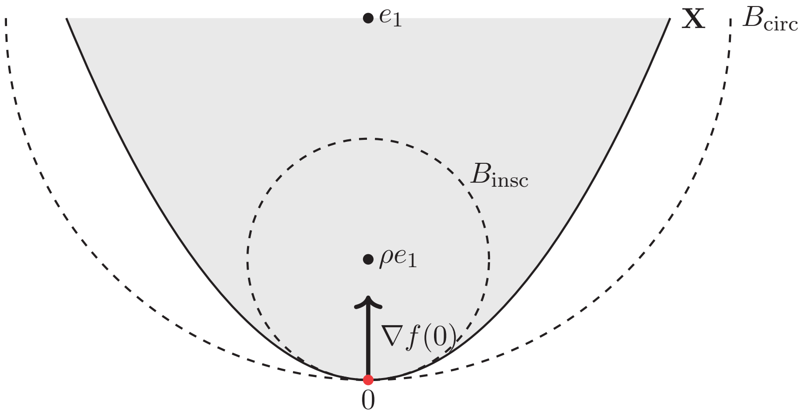

FIGURE 3.2. An overview of the situation in the second case of the proof of Theorem 3.23.

The argument is similar to the one used for the proof of Lemma 3.14. By (3.10), there is a constant $C > 0$ such that $\int_0^2 zq_r(z)w'(z)dz \leq \frac{C}{r^2}$ for all $r \in \mathbb{N}$. So we have:

$$\frac{C}{r^2} \geq \int_\rho^2 zq_r(z)w'(z)dz \geq \rho \int_\rho^2 q_r(z)w'(z)dz = \rho\left(1 - \int_0^\rho q_r(z)w'(z)dz\right),$$

which implies $\int_0^\rho q_r(z)w'(z)dz \geq 1 - \frac{C}{\rho r^2} \geq \frac{1}{2}$ for r large enough.

This concludes the proof of Theorem 3.23. □

3.5.2. Classification of ball-like sets. With Theorem 3.23 in mind, it is interesting to understand under which conditions a convex body $\mathbf{X}$ is ball-like. Under the assumption that $\mathbf{X}$ has a C^2-boundary, the well-known Rolling Ball Theorem (cf., e.g., [**Kou72**]) guarantees the existence of inscribed tangent balls.

THEOREM 3.24 (Rolling Ball Theorem). *Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a convex body with C^2- boundary. Then there exists $\varepsilon_{\text{insc}} > 0$ such that $\mathbf{X}$ has an inscribed tangent ball of radius $\varepsilon_{\text{insc}}$ for each $\mathbf{x} \in \partial\mathbf{X}$.*

Classifying the existence of circumscribed tangent balls is somewhat more involved. Certainly, we should assume that $\mathbf{X}$ is *strictly* convex, which means that its boundary should not contain any line segments. This assumption, however, is not sufficient. Instead we need the following stronger notion of *2-strict convexity* introduced in [**DH06**].

DEFINITION 3.25. *Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a convex body with C^2-boundary and let $\Psi \in C^2(\mathbb{R}^n, \mathbb{R})$ such that $\mathbf{X} = \Psi^{-1}((-\infty, 0])$ and $\partial\mathbf{X} = \Psi^{-1}(0)$. Assume $\nabla\Psi(a) \neq 0$ for all $a \in \partial\mathbf{X}$. The set $\mathbf{X}$ is said to be* 2-strictly convex *if the* *following holds:*

$$\mathbf{x}^T \nabla^2 \Psi(a)\mathbf{x} > 0 \quad \text{for all } \mathbf{x} \in T_a\mathbf{X} \setminus \{0\} \text{ and } a \in \partial\mathbf{X}.$$

In other words, the Hessian of Ψ at any boundary point should be positive definite, when restricted to the tangent space.

EXAMPLE 3.26. *Consider the unit ball for the ℓ_4-norm:*

$$\mathbf{X} = \{(\mathbf{x}_1, \mathbf{x}_2) : \Psi(\mathbf{x}_1, \mathbf{x}_2) := \mathbf{x}_1^4 + \mathbf{x}_2^4 \leq 1\} \subseteq \mathbb{R}^2.$$

Then, $\mathbf{X}$ is strictly convex, but not 2-strictly convex. Indeed, at any of the points $a = (0, \pm 1)$ and $(\pm 1, 0)$, the Hessian of Ψ is not positive definite on the tangent space. For instance, for $a = (0, -1)$, we have $\nabla\Psi(a) = (0, -4)$ and $\mathbf{x}^T\Psi^2(a)\mathbf{x} = 12\mathbf{x}_2^2$, which vanishes at $\mathbf{x} = (1, 0) \in T_a\mathbf{X}$. In fact, one can verify that $\mathbf{X}$ does not have a circumscribed tangent ball at any of the points $(0, \pm 1)$, $(\pm 1, 0)$.

It is shown in [**DH06**] that the set of 2-strictly convex bodies lies dense in the set of all convex bodies. For $\mathbf{X}$ with C^2-boundary, it turns out that 2-strict convexity is equivalent to the existence of circumscribed tangent balls at all boundary points.

THEOREM 3.27 ([**DS07**, Corollary 3.3]). *Let $\mathbf{X}$ be a convex body with C^2-boundary. Then $\mathbf{X}$ is 2-strictly convex if and only if there exists $\varepsilon_{\text{circ}} > 0$ such that $\mathbf{X}$ has a circumscribed tangent ball of radius $\varepsilon_{\text{circ}}$ at all boundary points $a \in \partial\mathbf{X}$.*

Combining Theorems 3.24 and 3.27 then gives a full classification of the ball-like convex bodies $\mathbf{X}$ with C^2-boundary.

COROLLARY 3.28. *Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a convex body with C^2-boundary. Then $\mathbf{X}$ is ball-like if and only if it is 2-strictly convex.*

3.5.3. A convex body without inscribed tangent balls. We now give an example of a convex body $\mathbf{X}$ which does not have inscribed tangent balls, going back to de Rham [**dR47**]. The idea is to construct a curve by starting with a polygon, and then successively 'cutting corners'. Let C_0 be the polygon in $\mathbb{R}^2$ with vertices $(-1, -1), (1, -1), (1, 1)$ and $(-1, 1)$, i.e., a square. For $k \geq 1$, we obtain C_k by subdividing each edge of C_{k-1} into three equal parts and taking the convex hull of the resulting subdivision points (see Figure 3.3). We then let C be the limiting curve obtained by letting k tend to ∞. Then, C is a continuously differentiable, convex curve (see [**dB87**] for details). It is not, however, C^2 everywhere. We indicate below some point where no inscribed tangent ball exists for the convex body with boundary C.

Consider the point $m = (0, -1) \in C$, which is an element of C_k for all k. Fix $k \geq 1$. If we walk anti-clockwise along C_k starting at m, the first corner point encountered is $s_k = (1/3^k, -1)$, the slope of the edge starting at s_k is $l_k = 1/k$ and its end point is

$$e_k = \big((2k + 1)/3^k, 2/3^k - 1\big).$$

Now suppose that there exists an inscribed tangent ball $B_\varepsilon(c)$ at the point m. Then, $\varepsilon > 0$, $c = (0, \varepsilon - 1)$ and any point $(x, y) \in C$ lies outside of the ball $B_\varepsilon(c)$, so that

$$x^2 + (y+1)^2 - 2\varepsilon(y+1) \geq 0 \quad \text{for all } (x, y) \in C.$$

As C is contained in the polygonal region delimited by any C_k, also $e_k \notin B_\varepsilon(c)$ and thus $\left(\frac{2k+1}{3^k}\right)^2 + \left(\frac{2}{3^k}\right)^2 - \frac{4\varepsilon}{3^k} \geq 0$. Letting $k \to \infty$, we get $\varepsilon = 0$, a contradiction.

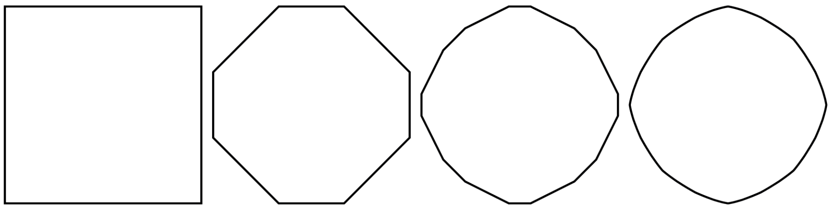

FIGURE 3.3. From left to right: the curve C_k for $k = 0, 1, 2, 8$.

3.6. The simplex

We now consider a full-dimensional simplex $\Delta^n = \operatorname{conv}(\{v_0, v_1, v_2, \ldots, v_n\})$ in $\mathbb{R}^n$, equipped with the Lebesgue measure. We show the following.

THEOREM 3.29. *Let* $\mathbf{X} = \Delta^n$ *be a simplex, equipped with the Lebesgue measure. Then*

$$\operatorname{Error}(f; \Delta^n)_r = O\left(\frac{1}{r^2}\right).$$

PROOF. Let $\mathbf{x}^* \in \Delta^n$ be a global minimizer of f on Δ^n. The idea is to apply an affine transformation ϕ to Δ^n whose image $\phi(\Delta^n)$ is locally similar to $[0,1]^n$ at the global minimizer $\phi(\mathbf{x}^*)$ of $g := f \circ \phi^{-1}$, after which we may 'transport' the $O(1/r^2)$ rate from the hypercube to the simplex.

Let $F := \operatorname{conv}(v_1, v_2, \ldots, v_n)$ be the facet of Δ^n which does not contain v_0. By reindexing, we may assume w.l.o.g. that $\mathbf{x}^* \notin F$. Consider the map ϕ determined by $\phi(v_0) = 0$ and $\phi(v_i) = e_i$ for all $i \in [n]$, where e_i is the i-th standard basis vector of $\mathbb{R}^n$. See Figure 3.4. Clearly, ϕ is nonsingular, and $\phi(\Delta^n) \subseteq [0,1]^n$.

LEMMA 3.30. *We have* $\phi(\Delta^n) \subseteq_{\phi(\mathbf{x})} [0,1]^n$ *for all* $\mathbf{x} \in \Delta^n \setminus F$.

PROOF. By definition of F, we have

$$\Delta^n \setminus F = \left\{ \sum_{i=0}^n \lambda_i v_i : \sum_{i=1}^n \lambda_i < 1, \lambda \geq 0 \right\},$$

and so

$$\phi(\Delta^n \setminus F) = \left\{ y \in [0,1]^n : \sum_{i=1}^n y_i < 1 \right\},$$

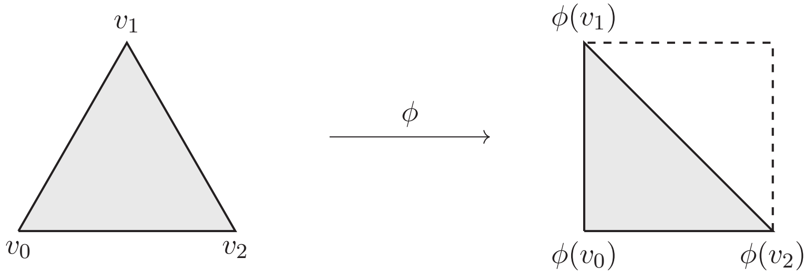

FIGURE 3.4. The map ϕ from the proof of Theorem 3.29 for $n = 2$

which is an open subset of $[0,1]^n$. But this means that for each $\mathbf{y} = \phi(\mathbf{x}) \in \phi(\Delta^n \setminus F)$ there exists $\varepsilon > 0$ such that

$$B^n_\varepsilon(\mathbf{y}) \cap [0,1]^n \subseteq B^n_\varepsilon(\mathbf{y}) \cap \phi(\Delta^n \setminus F),$$

which concludes the proof of the lemma. □

The above lemma tells us in particular that $\phi(\Delta^n) \subseteq_{\phi(\mathbf{x}^*)} [0,1]^n$. We now apply Corollary 3.15 with $\mathbf{X} = \phi(\Delta^n)$, $\widehat{\mathbf{X}} = [0,1]^n$ and weight functions $w = \widehat{w} = 1$ on $\mathbf{X}, \widehat{\mathbf{X}}$, respectively. This yields a polynomial upper estimator h of g on $[0,1]^n$ having:

$$\text{Error}(g; \phi(\Delta^n))_r \leq 2\text{Error}(h; [0,1]^n)_r = O(1/r^2),$$

for $r \in \mathbb{N}$ large enough, using Theorem 3.17 for the right most equality. It remains to apply Lemma 3.2 to obtain:

$$\text{Error}(f; \Delta^n)_r = \text{Error}(g; \phi(\Delta^n))_r = O(1/r^2),$$

which concludes the proof of Theorem 3.29. □

3.7. Discussion

We have shown that the measure-based upper bounds $\text{ub}(f, \mathbf{X}, \mu)_r$ converge to the global minimum $f_{\min}$ of f on $\mathbf{X}$ at a rate in $O(1/r^2)$, for a class of special convex bodies $\mathbf{X}$, w.r.t. natural references measures μ. We reiterate that this rate is best-possible in general. In light of the results of this chapter, the main open question is whether the convergence rate in $O(1/r^2)$ can be extended to *all* convex bodies.

In particular, it is interesting to determine the exact rate of convergence for polytopes. We could so far only deal with hypercubes and simplices. The main tool we used was the 'local similarity' of the simplex with the hypercube. For a general polytope $\mathbf{X}$, if the minimum is attained at a point lying in the interior of $\mathbf{X}$ or of one of its facets, then we can still apply the 'local similarity' tool (and deduce the $O(1/r^2)$ rate). However, at other points (like its vertices) $\mathbf{X}$ is in general not locally similar to the hypercube, so another proof technique seems needed. A possible strategy could be splitting $\mathbf{X}$ into simplices and using the known convergence rate for the simplex containing a global minimizer; however, a difficulty there is keeping track of the distribution of mass of an optimal sum-of-squares on the other simplices.

As we shall see in the next chapter, it is at least possible to show a rate in $O(\log^2(r)/r^2)$ for all convex bodies (in fact, for all compact semialgebraic sets with dense interior). This rate is thus only a log-factor away from the rates shown in this chapter.

CHAPTER 4

Convergence analysis for measure-based bounds II

Everything is more complicated than you think. You only see a tenth of what is true.

From Synecdoche, New York

This chapter is based on my joint works [**SL20, SL21a**] *with Monique Laurent.*

Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact set, and let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree d to be minimized over $\mathbf{X}$. Recall from Chapter 2 the measure-based upper bound on the global minimum $f_{\min}$ of f:

$$\text{ub}(f, \mathbf{X})_r := \inf_{q \in \Sigma[\mathbf{x}]_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mathbf{x} : \int_{\mathbf{X}} q(\mathbf{x}) d\mathbf{x} = 1 \right\}. \tag{4.1}$$

Recall also the push-forward bound:

$$\text{ub}(f, \mathbf{X})_r^{\text{pf}} := \inf_{u \in \Sigma[x]_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) u(f(\mathbf{x})) d\mathbf{x} : \int_{\mathbf{X}} u(f(\mathbf{x})) d\mathbf{x} = 1 \right\}, \tag{4.2}$$

which is obtained from (4.1) by restricting the minimization to sums of squares q of the form $q = u \circ f$ with $u \in \Sigma[x]$. It therefore satisfies:

$$\text{ub}(f, \mathbf{X})_r^{\text{pf}} \geq \text{ub}(f, \mathbf{X})_{dr} \geq f_{\min}. \tag{4.3}$$

Consider the following geometric assumption on the set $\mathbf{X}$.

Assumption 4.1. *There exist positive constants $\varepsilon_{\mathbf{X}}, \eta_{\mathbf{X}} > 0$ and $N \geq n$, such that, for all $\mathbf{x} \in \mathbf{X}$ and $0 < \delta \leq \varepsilon_{\mathbf{X}}$, we have:*

$$\text{vol}(\mathbf{X} \cap B_\delta^n(\mathbf{x})) \geq \eta_{\mathbf{X}} \delta^N \text{vol}(B^n). \tag{4.4}$$

Here, $B_\delta^n(\mathbf{x})$ is the Euclidean ball centered at $\mathbf{x}$ with radius δ and $B^n = B_1^n(0)$.

A slightly stronger version of Assumption 4.1 (requiring $N = n$) was introduced in [**dKLS17**], where it was used to give the first error analysis in $O(1/\sqrt{r})$ for the bounds $\text{ub}(f, \mathbf{X})_r$. The condition of [**dKLS17**] is satisfied, e.g., when $\mathbf{X}$ is a convex body, or more generally when $\mathbf{X}$ satisfies an interior cone condition, or when $\mathbf{X}$ is star-shaped with respect to a ball (see also [**dKLS17**] for a more complete discussion). The weaker condition of Assumption 4.1 is satisfied additionally by compact semialgebraic sets that have a dense interior, which permits in particular that $\mathbf{X}$ has certain types of cusps. We discuss Assumption 4.1 in more detail in Section 4.4 below.

Outline. In this chapter, we show a convergence rate in $O(\log^2 r/r^2)$ for the measure-based hierarchy $\text{ub}(f, \mathbf{X})_r$ when applied to optimization over $\mathbf{X}$ satisfying this assumption. In fact, we show that this rate holds already for the (weaker) push-forward bound $\text{ub}(f, \mathbf{X})_r^{\text{pf}}$. This rate is only a log-factor away from the best-possible rate in $O(1/r^2)$, see Section 2.3.

THEOREM 4.2. *Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact set satisfying Assumption 4.1 and let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial. Then we have that:*

$$\text{ub}(f, \mathbf{X})_r^{\text{pf}} - f_{\min} = O(\log^2 r/r^2).$$

As a direct consequence, we also have $\text{ub}(f, \mathbf{X})_r - f_{\min} = O(\log^2 r/r^2)$.

Theorem 4.2 improves upon the earlier results of [**dKLS17, dKL18**] in three ways. First, it applies to a (much) broader class of compact sets $\mathbf{X}$. Second, it shows a better convergence rate in $O(\log^2 r/r^2)$, as opposed to $O(1/\sqrt{r})$ and $O(1/r)$, respectively. Finally, it applies not only to the (regular) measure-based bounds $\text{ub}(f, \mathbf{X})_r$, but also to the more economical push-forward bounds $\text{ub}(f, \mathbf{X})_r^{\text{pf}}$.

The idea of our proof is to exhibit for each $r \in \mathbb{N}$ an explicit sum-of-squares density on $\mathbf{X}$ of the form $u \circ f$, $u \in \Sigma[x]_{2r}$ for which $\int_{\mathbf{X}} f(\mathbf{x})u(f(\mathbf{x}))d\mathbf{x}$ is small. That is, to construct a feasible solution to (4.2) with suitable objective value. This method of proof is in the same spirit as the one of the earlier works [**dKLS17, dKL18**], see also Chapter 2. Our construction differs from the one in these works in two ways. First, it makes use of so-called *needle polynomials* to approximate the Dirac function on $[-1, 1]$ or $[0, 1]$. This yields a better approximation than the ones of [**dKLS17, dKL18**], which are based on truncated Taylor expansions of a Gaussian function. Second, it makes use of the push-forward bounds to reduce the analysis to a *univariate* setting, which greatly simplifies the proof.

As a small additional result, we show in Section 4.3 that there exists a class of polynomials f on the interval $\mathbf{X} = [-1, 1]$ for which the regular measure-based bounds converge to $f_{\min}$ much much faster than the (weaker) push-forward bounds. This result contrasts the (somewhat surprising) fact that the general worst-case convergence rate for these two types of bounds is quite similar; namely within a logarithmic factor of $O(1/r^2)$.

4.1. Needle polynomials

In this section, we cover some preliminaries needed to prove our main result; in particular we introduce the *needle polynomials* central to our construction and derive some of their technical properties.

We begin by recalling some basic properties of the Chebyshev polynomials, see also Chapter 1. The Chebyshev polynomials $\mathcal{C}_r \in \mathbb{R}[x]_r$ may be defined by the following explicit expression:

$$\mathcal{C}_r(x) = \begin{cases} \cos(r \arccos x) & \text{for } |x| \leq 1, \\ \frac{1}{2}(t + \sqrt{x^2-1})^r + \frac{1}{2}(x - \sqrt{x^2-1})^r & \text{for } |x| \geq 1. \end{cases} \tag{4.5}$$

From this definition, it can be seen that $|\mathcal{C}_r(x)| \leq 1$ on the interval $[-1, 1]$, and that $\mathcal{C}_r(x)$ is nonnegative and monotone nondecreasing on $[1, \infty)$. The Chebyshev polynomials form an orthogonal basis of $\mathbb{R}[x]$ with respect to the Chebyshev measure (with weight function $(1-x^2)^{-1/2}$) on $[-1, 1]$ and they are used extensively in approximation theory. For instance, they are the polynomials attaining equality in the Markov brother's inequality on $[-1, 1]$, recalled below.

Lemma 4.3 (Markov Brothers' Inequality; see, e.g., [**Sha04**]). *Let $p \in \mathbb{R}[x]$ be a univariate polynomial of degree at most r. Then, for any scalars $a < b$, we have:*

$$\max_{x\in[a,b]} |p'(x)| \leq \frac{2r^2}{b-a} \cdot \max_{x\in[a,b]} |p(x)|.$$

Kroó and Swetits [**KS92**] use the Chebyshev polynomials to construct the so-called (univariate) *needle polynomials.*

Definition 4.4. *For $r \in \mathbb{N}, h \in (0, 1)$, we define the* needle polynomial *$\nu_r^h \in \mathbb{R}[x]_{4r}$ by*

$$\nu_r^h(x) = \frac{\mathcal{C}_r^2(1 + h^2 - x^2)}{\mathcal{C}_r^2(1 + h^2)}.$$

Additionally, we define the $\frac{1}{2}$-needle polynomial *$\widehat{\nu}_r^h \in \mathbb{R}[x]_{4r}$ by*

$$\widehat{\nu}_r^h(x) = \mathcal{C}_{2r}^2\left(\frac{2 + h - 2x}{2 - h}\right) \cdot \mathcal{C}_{2r}^{-2}\left(\frac{2 + h}{2 - h}\right).$$

By construction, the needle polynomials ν_r^h and $\widehat{\nu}_r^h$ are squares and have degree $4r$. They approximate well the Dirac delta function at 0 on $[-1, 1]$ and $[0, 1]$, respectively. In [**Sen90**], a construction similar to the needles presented here is used to obtain the best polynomial approximation of the Dirac delta in terms of the Hausdorff distance. In our proof of Theorem 4.2 below, we will only need the $\frac{1}{2}$-needle polynomials. For completeness, however, we will cover both variants here. As we discuss briefly at the end of this Chapter, the (regular) needle polynomials were actually used to analyse the measure-based hierarchy in [**SL20**], but this analysis has since been made obsolete by the results presented in this chapter.

The needle polynomials satisfy the following bounds (see Figure 4.1 and Figure 4.2 for an illustration).

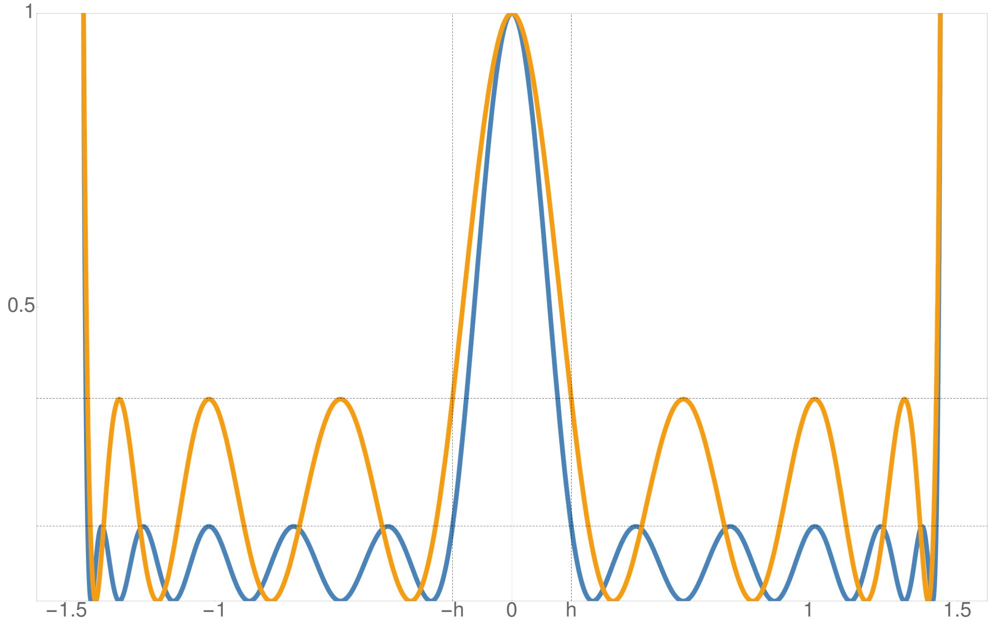

FIGURE 4.1. The needle polynomials ν_4^h (orange), ν_6^h (blue) for $h = 1/5$.

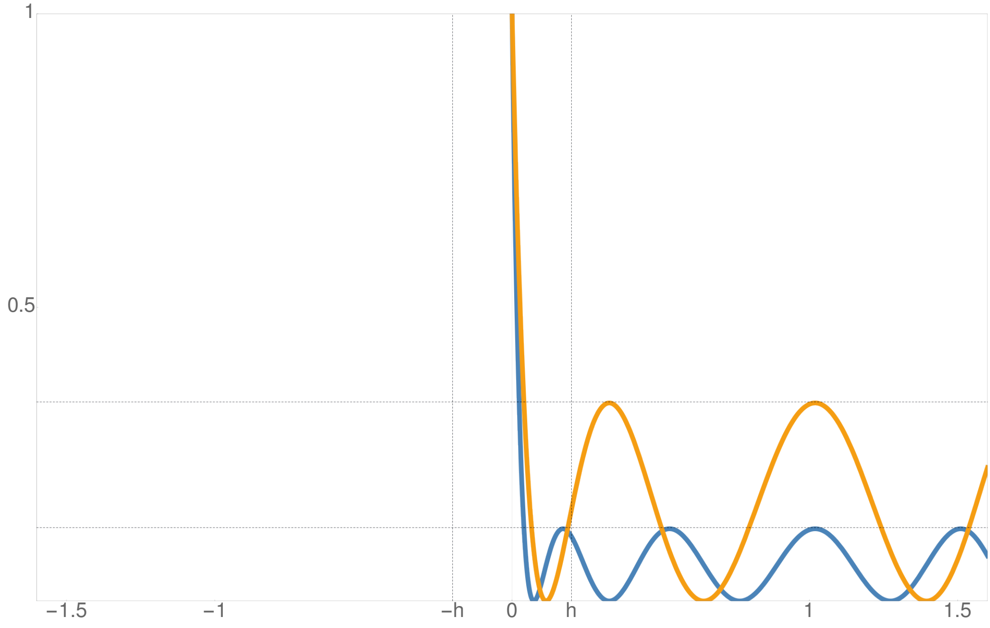

FIGURE 4.2. The $\frac{1}{2}$-needle polynomials $\widehat{\nu}_4^{h^2}$ (orange), $\widehat{\nu}_6^{h^2}$ (blue) for $h = 1/5$.

THEOREM 4.5 (cf. [**KS92, KL12, Kro15**]). *For any $r \in \mathbb{N}$ and $h \in (0,1)$, the following properties hold for the polynomials ν_r^h and $\widehat{\nu}_r^h$:*

$$\nu_r^h(0) = 1, \tag{4.6}$$

$$0 \le \nu_r^h(x) \le 1 \qquad \text{for } x \in [-1,1], \tag{4.7}$$

$$\nu_r^h(x) \le 4e^{-\frac{1}{2}rh} \qquad \text{for } x \in [-1,1] \text{ with } |x| \ge h, \tag{4.8}$$

$$\widehat{\nu}_r^h(0) = 1, \tag{4.9}$$

$$0 \le \widehat{\nu}_r^h(x) \le 1 \qquad \text{for } x \in [0,1], \tag{4.10}$$

$$\widehat{\nu}_r^h(x) \le 4e^{-\frac{1}{2}r\sqrt{h}} \qquad \text{for } x \in [0,1] \text{ with } x \ge h. \tag{4.11}$$

As this result plays a central role in our treatment we give a short proof, following the argument given in [**LP19**]. We need the following lemma.

LEMMA 4.6. *For any $r \in \mathbb{N}$, $x \in [0,1)$, we have:*

$$\mathcal{C}_r(1+x) \ge \frac{1}{2}e^{r\sqrt{x}\log(1+\sqrt{2})} \ge \frac{1}{2}e^{\frac{1}{4}r\sqrt{x}}.$$

PROOF. Using the explicit expression (4.5) for $\mathcal{C}_r$, we have:

$$\begin{aligned} 2\mathcal{C}_r(1+x) &\ge \left(1 + x + \sqrt{(1+x)^2 - 1}\right)^r = (1 + x + \sqrt{2x + x^2})^r \\ &\ge (1 + \sqrt{2x})^r = e^{r\log(1+\sqrt{2}\cdot\sqrt{x})}. \end{aligned}$$

By concavity of the logarithm, we have:

$$\begin{aligned} \log(1 + \sqrt{2}\sqrt{x}) &= \log(\sqrt{x}\cdot(1+\sqrt{2}) + (1-\sqrt{x})\cdot 1) \\ &\ge \sqrt{x}\cdot\log(1+\sqrt{2}) + (1-\sqrt{x})\log(1) \\ &= \sqrt{x}\cdot\log(1+\sqrt{2}) \ge \frac{1}{4}\sqrt{x}, \end{aligned}$$

and so, using the above lower bound on $\mathcal{C}_r(1+x)$, we obtain:

$$\mathcal{C}_r(1+x) \ge \frac{1}{2}e^{r\sqrt{x}\log(1+\sqrt{2})} \ge \frac{1}{2}e^{\frac{1}{4}r\sqrt{x}}.$$

□

PROOF OF THEOREM 4.5. Properties (4.6), (4.9) are clear. We first check (4.7)-(4.8). If $|x| \le h$ then $1 + h^2 \ge 1 + h^2 - x^2 \ge 1$, giving $\nu_r^h(x) \le \nu_r^h(0) = 1$ by monotonicity of $\mathcal{C}_r(x)$ on $[1,\infty)$. Assume now $h \le |x| \le 1$. Then $\mathcal{C}_r^2(1 + h^2 - t^2) \le 1$ as $1 + h^2 - x^2 \in [-1,1]$, and $\mathcal{C}_r^2(1+h^2) \ge 1$ (again by monotonicity), which implies $\nu_r^h(x) \le 1$. In addition, since $\mathcal{C}_r(1+h^2) \ge \frac{1}{2}e^{\frac{1}{4}rh}$ by Lemma 4.6, we obtain $\nu_r^h(x) \le \mathcal{C}_r^{-2}(1+h^2) \le 4e^{-\frac{1}{2}rh}$.

We now check (4.10)-(4.11). If $x \in [0,h]$ then $\widehat{\nu}_r^h(x) \le \widehat{\nu}_r^h(0) = 1$ follows by monotonicity of $\mathcal{C}_{2r}(x)$ on $[1,\infty)$. Assume now $h \le x \le 1$. Then,

$\frac{2+h-2x}{2-h} \in [-1,1]$ and thus $\mathcal{C}_{2r}^2\big(\frac{2+h-2x}{2-h}\big) \leq 1$. On the other hand, we have $\mathcal{C}_{2r}^2\big(\frac{2+h}{2-h}\big) \geq 1$, which gives $\widehat{\nu}_r^h(x) \leq 1$. In addition, as $\frac{2+h}{2-h} \geq 1+h \geq 1$, using again monotonicity of $\mathcal{C}_{2r}$ and Lemma 4.6, we get $\mathcal{C}_{2r}^2\big(\frac{2+h}{2-h}\big) \geq \mathcal{C}_{2r}^2(1+h) \geq \frac{1}{4}e^{\frac{1}{2}r\sqrt{h}}$, which implies (4.11).

□

The following lemma gives a simple lower estimator on $[0,1]$ for a nonnegative polynomial p of given degree with $p(0)=1$. This lower estimator allows us to lower bound the value of the $\frac{1}{2}$-needle polynomials on small intervals $[0,h]$, $h>0$. Such a lower bound will be useful in the proof of Theorem 4.2 below.

LEMMA 4.7. *Let $p \in \mathbb{R}[x]_r$ be a polynomial, which is nonnegative over $\mathbb{R}_{\geq 0}$ and satisfies $p(0)=1$, $p(x) \leq 1$ for all $x \in [0,1]$. Let $\Lambda_r : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be defined by:*

$$\Lambda_r(x) = \begin{cases} 1-2r^2x & \text{if } x \leq \frac{1}{2r^2}, \\ 0 & \text{otherwise.} \end{cases}$$

Then $\Lambda_r(x) \leq p(x)$ for all $x \in \mathbb{R}_{\geq 0}$.

PROOF. Suppose not. Then there exists $s \in \mathbb{R}_{\geq 0}$ such that $\Lambda_r(s) > p(s)$. As $p \geq 0$ on $\mathbb{R}_{\geq 0}$, $p(0)=1$ and $\Lambda_r(x)=0$ for $x \geq \frac{1}{2r^2}$, we have $0 < s < \frac{1}{2r^2}$. We find that $p(s)-p(0) < \Lambda_r(s)-1 = -2r^2 s$. Now, by the mean value theorem, there exists an element $z \in (0,s)$ such that $p'(z) = \frac{p(s)-p(0)}{s} < \frac{-2r^2 s}{s} = -2r^2$. But this is in contradiction with Lemma 4.3, which implies that $\max_{x\in[0,1]} |p'(x)| \leq 2r^2$. □

COROLLARY 4.8. *Let $h \in (0,1)$, and let $\nu_r^h, \widehat{\nu}_r^h$ as above. Then $\Lambda_{4r}(x) \leq \nu_r^h(x) = \nu_r^h(-x)$ and $\Lambda_{4r}(x) \leq \widehat{\nu}_r^h(x)$ for all $x \in [0,1]$. In particular, we thus have that:*

$$\widehat{\nu}_r^h(x) \geq 1-32r^2x \geq \frac{1}{2} \quad \textit{for all } x \in [0, \frac{1}{64r^2}].$$

4.2. Proof of the main result

We show the following restatement of Theorem 4.2.

THEOREM 4.9. *Assume $\mathbf{X}$ is connected compact and satisfies the above geometric condition (4.4). Then there exists a constant C (depending only on n, the Lipschitz constant of f and $\mathbf{X}$) such that:*

$$\mathrm{ub}(f,\mathbf{X})_r^{\mathrm{pf}} - f_{\min} \leq C\frac{\log^2 r}{r^2}(f_{\max}-f_{\min}) \quad \textit{for all large } r.$$

The rest of this section is devoted to the proof of Theorem 4.9. We will make the following assumptions in order to simplify notation in our arguments. Let $\mathbf{x}^*$ be a global minimizer of f in $\mathbf{X}$. After applying a suitable translation (replacing $\mathbf{X}$ by $\mathbf{X}-\mathbf{x}^*$ and the polynomial f by the polynomial $\mathbf{x} \mapsto f(\mathbf{x}-\mathbf{x}^*)$), we may assume that $\mathbf{x}^* = 0$, that is, we may assume that the global minimum of f over $\mathbf{X}$ is attained at the origin. Furthermore, it suffices to work with the rescaled polynomial:

$$F(x) := \frac{f(\mathbf{x}) - f_{\min}}{f_{\max} - f_{\min}},$$

which satisfies $F(\mathbf{X}) = [0, 1]$, with $F_{\min} = 0$ and $F_{\max} = 1$. Indeed, one can easily check that:

$$\mathrm{ub}(f, \mathbf{X})_r^{\mathrm{pf}} - f_{\min} \leq (f_{\max} - f_{\min})\mathrm{ub}(F, \mathbf{X})_r^{\mathrm{pf}}.$$

For clarity, we will denote by μ the Lebesgue measure on $\mathbb{R}^n$. Recall that the push-forward measure μ_F is defined by:

$$\mu_F(B) = \mu(F^{-1}(B)) \quad (B \subseteq \mathbb{R} \text{ Borel}). \tag{4.12}$$

Because of the scaling, the support of the μ_F is equal to $[0, 1]$, and so:

$$\begin{aligned} \mathrm{ub}(F, \mathbf{X})_r^{\mathrm{pf}} &= \inf_{s \in \Sigma[x]_{2r}} \Big\{ \int_0^1 x s(x) d\mu_F(x) : \int_0^1 s(x) d\mu_F(x) = 1 \Big\} \\ &= \inf_{s \in \Sigma[x]_{2r}} \Big\{ \int_{\mathbf{X}} F(\mathbf{x}) s(F(\mathbf{x})) d\mu(\mathbf{x}) : \int_{\mathbf{X}} s(F(\mathbf{x})) d\mu(\mathbf{x}) = 1 \Big\}. \end{aligned} \tag{4.13}$$

In order to analyze the bound $\mathrm{ub}(F, \mathbf{X})_r^{\mathrm{pf}}$, we construct a univariate sum-of-squares polynomial s which approximates well the Dirac delta centered at the origin on the interval $[0, 1]$, making use of the $\frac{1}{2}$-needle polynomials introduced above. Indeed, we consider the sum-of-squares polynomial $s(x) := C\widehat{\nu}_r^h(x)$, where $h \in (0, 1)$ will be chosen later, and C is chosen so that s is a density on $[0, 1]$ with respect to the measure μ_F. That is:

$$C = \left(\int_{\mathbf{X}} \widehat{\nu}_r^h(F(\mathbf{x})) d\mu(\mathbf{x}) \right)^{-1}.$$

Therefore, s is a feasible solution to (4.13), and so we obtain:

$$\mathrm{ub}(F, \mathbf{X})_r^{\mathrm{pf}} \leq \int_{\mathbf{X}} F(\mathbf{x}) s(F(\mathbf{x})) d\mu(\mathbf{x}) = \frac{\int_{\mathbf{X}} F(\mathbf{x}) \widehat{\nu}_r^h(F(\mathbf{x})) d\mu(\mathbf{x})}{\int_{\mathbf{X}} \widehat{\nu}_r^h(F(\mathbf{x})) d\mu(\mathbf{x})}.$$

Our goal is thus to show that:

$$\mathrm{ratio} := \frac{\int_{\mathbf{X}} F(\mathbf{x}) \widehat{\nu}_r^h(F(\mathbf{x})) d\mu(\mathbf{x})}{\int_{\mathbf{X}} \widehat{\nu}_r^h(F(\mathbf{x})) d\mu(\mathbf{x})} = O\left(\frac{\log^2 r}{r^2} \right). \tag{4.14}$$

Define the set

$$\mathbf{X}_h = \{\mathbf{x} \in \mathbf{X} : F(\mathbf{x}) \leq h\}.$$

We first work out the numerator of (4.14), which we split into two terms, depending whether we integrate on $\mathbf{X}_h$ or on its complement:

$$\begin{aligned} \int_{\mathbf{X}} F(\mathbf{x}) \widehat{\nu}_r^h(F(\mathbf{x})) d\mu(\mathbf{x}) &= \int_{\mathbf{X}_h} F(\mathbf{x}) \widehat{\nu}_r^h(F(\mathbf{x})) d\mu(\mathbf{x}) + \int_{\mathbf{X} \setminus \mathbf{X}_h} F(\mathbf{x}) \widehat{\nu}_r^h(F(\mathbf{x})) d\mu(\mathbf{x}) \\ &\leq h \int_{\mathbf{X}_h} \widehat{\nu}_r^h(F(\mathbf{x})) d\mu(\mathbf{x}) + \int_{\mathbf{X} \setminus \mathbf{X}_h} \widehat{\nu}_r^h(F(\mathbf{x})) d\mu(\mathbf{x}). \end{aligned}$$

Here we have upper bounded $F(\mathbf{x})$ by h on $\mathbf{X}_h$ and by 1 on $\mathbf{X} \setminus \mathbf{X}_h$. On the other hand, we can lower bound the denominator in (4.14) as follows:

$$\int_{\mathbf{X}} \widehat{\nu}_r^h(F(\mathbf{x}))d\mu(\mathbf{x}) \geq \int_{\mathbf{X}_h} \widehat{\nu}_r^h(F(\mathbf{x}))d\mu(\mathbf{x}).$$

Combining the above two inequalities on the numerator and denominator we get:

$$\text{ratio} \leq h + \frac{\int_{\mathbf{X}\setminus\mathbf{X}_h} \widehat{\nu}_r^h(F(\mathbf{x}))d\mu(\mathbf{x})}{\int_{\mathbf{X}_h} \widehat{\nu}_r^h(F(\mathbf{x}))d\mu(\mathbf{x})}.$$

Thus we only need to upper bound the second term above. We first work on the numerator. For any $x \in \mathbf{X} \setminus \mathbf{X}_h$ we have $F(\mathbf{x}) > h$ and thus, using Theorem 4.5, we get $\widehat{\nu}_r^h(F(\mathbf{x})) \leq 4e^{-\frac{1}{2}r\sqrt{h}}$. This implies that:

$$\int_{\mathbf{X}\setminus\mathbf{X}_h} \widehat{\nu}_r^h(F(\mathbf{x}))d\mu(\mathbf{x}) \leq 4e^{-\frac{1}{2}r\sqrt{h}}\mu(\mathbf{X}).$$

Next, we bound the denominator. Corollary 4.8 tells us that:

$$\widehat{\nu}_r^h(x) \geq 1 - 32r^2x \geq \frac{1}{2} \quad \text{for all } x \in [0, \frac{1}{64r^2}].$$

Set $\rho = \frac{1}{64r^2}$. We will later choose $h \geq \rho$, so that $\mathbf{X}_h \supseteq \mathbf{X}_\rho := \{\mathbf{x} \in \mathbf{X} : F(\mathbf{x}) \leq \rho\}$ and $\widehat{\nu}_r^h(F(\mathbf{x})) \geq \frac{1}{2}$ for all $\mathbf{x} \in \mathbf{X}_\rho$. As $\mathbf{X}$ is compact, there exists a Lipschitz constant $C_F > 0$ such that:

$$F(\mathbf{x}) \leq C_F\|\mathbf{x}\| \quad \text{for all } \mathbf{x} \in \mathbf{X}.$$

Note that $\mathbf{X} \cap B^n_{\rho/C_F} \subseteq \mathbf{X}_\rho$. By the geometric assumption (4.4), we have:

$$\mu(\mathbf{X} \cap B^n_{\rho/C_F}) \geq \eta_{\mathbf{X}} \left(\frac{\rho}{C_F}\right)^N \mu(B^n)$$

for all r large enough such that $\rho/C_F \leq \varepsilon_{\mathbf{X}}$. We can then lower bound the denominator as follows:

$$\begin{aligned}\int_{\mathbf{X}_h} \widehat{\nu}_r^h(F(\mathbf{x}))d\mu(\mathbf{x}) &\geq \int_{\mathbf{X}_\rho} \widehat{\nu}_r^h(F(\mathbf{x}))d\mu(\mathbf{x}) \geq \frac{1}{2}\mu(\mathbf{X}_\rho) \\ &\geq \frac{1}{2}\mu(\mathbf{X} \cap B^n_{\rho/C_F}) \geq \frac{1}{2}\eta_{\mathbf{X}} \left(\frac{\rho}{C_F}\right)^N \cdot \mu(B^n).\end{aligned}$$

Combining the above inequalities, we obtain:

$$\text{ratio} \leq h + \frac{e^{-\frac{1}{2}r\sqrt{h}}}{\rho^N} \cdot \frac{8 \cdot \mu(\mathbf{X})(C_F)^N}{\eta_{\mathbf{X}}\mu(B^n)}.$$

If we now select $h = \left(4(N+1)\frac{\log r}{r}\right)^2$, we have $h \geq \rho$ and a straightforward computation shows that:

$$\text{ratio} = O\left(\frac{\log^2 r}{r^2}\right).$$

Here, the constant in the big O depends on n, N, C_F, $\eta_{\mathbf{X}}$ and $\mu(\mathbf{X})$. This concludes the proof of Theorem 4.9.

4.3. Separation for a special class of polynomials

In this section we consider in more detail the behaviour of the bounds $\text{ub}(f, \mathbf{X})_r$ and $\text{ub}(f, \mathbf{X})_r^{\text{pf}}$ for the class of polynomials $f(x) = x^{2k}$ (with $k \geq 1$ integer) on the interval $\mathbf{X} = [-1, 1]$. Then $f([-1, 1]) = [0, 1]$ and, by applying (4.3) to the polynomial $f(x) = x^{2k}$, we have the following inequality:

$$0 \leq \text{ub}(f, \mathbf{X})_{2rk} \leq \text{ub}(f, \mathbf{X})_r^{\text{pf}} \quad \text{for any } r \geq 1.$$

Note that for any $i \leq 2k - 1$, the ith derivative of f vanishes at its global minimizer 0 on $[-1, 1]$. As we discuss a bit more at the end of this chapter, we may thus use [**SL20**, Theorem 14] to find that $\text{ub}(f, \mathbf{X})_{2rk} = O(\log^{2k} r/r^{2k})$. On the other hand, the convergence rate in $O(\log^2 r/r^2)$ for $\text{ub}(f, \mathbf{X})_r^{\text{pf}}$ shown in Theorem 4.2 is optimal up to the log-factor. Indeed, we will show here a lower bound on $\text{ub}(f, \mathbf{X})_r^{\text{pf}}$ in $\Omega(1/r^2)$.

Theorem 4.10. *Let $\mathbf{X} = [-1, 1]$ and let $f(x) = x^{2k}$ for $k \geq 1$ integer. Then we have $\text{ub}(f, \mathbf{X})_r^{\text{pf}} = \Omega(1/r^2)$.*

For $k \in \mathbb{N}$, let $\mu_k := \mu_f$ denote the push-forward measure (cf. (4.12)) of the Lebesgue measure on $[-1, 1]$ by the function $f(x) = x^{2k}$, and let $\{p_{k,i}(x) : i \in \mathbb{N}\} \subseteq \mathbb{R}[x]$ denote the family of orthogonal polynomials that provide an orthonormal basis for $\mathbb{R}[x]$ w.r.t. the inner product $\langle \cdot, \cdot \rangle_{\mu_k}$ (see Chapter 1). Then, as shown in [**dKL20b**] and as explained in Chapter 2, the parameter $\text{ub}(f, \mathbf{X})_r^{\text{pf}}$ is equal to the smallest root of the polynomial $p_{k,r+1}(x)$. As it turns out, the push-forward measure μ_k here is of Jacobi type. Hence, we have information about the corresponding orthogonal polynomials $p_{k,i}$, whose extremal roots are well understood. Let us recall from Chapter 1 the classical Jacobi polynomials.

Lemma 4.11. *Let $a, b > -1$. Consider the weight function $w_{a,b}(x) = (1-x)^a(1+x)^b$ on the interval $[-1, 1]$ and let $\{\mathcal{J}_k^{(a,b)}(x) : i \in \mathbb{N}\}$ be the corresponding family of orthogonal polynomials. Then $\mathcal{J}_i^{(a,b)}$ is the degree i Jacobi polynomial (with parameters a, b), and its smallest root $\xi_i^{a,b}$ satisfies:*

$$\xi_i^{a,b} = -1 + \Theta(1/i^2). \tag{4.15}$$

LEMMA 4.12. *For any integrable function g on $[-1,1]$ we have the identity:*

$$\int_{-1}^{1} g(x^{2k})dx = \frac{1}{k}\int_{0}^{1} g(x)x^{-1+1/2k}dx.$$

Hence, the push-forward measure μ_k is given by $d\mu_k(x) := \frac{1}{k}x^{-1+\frac{1}{2k}}dx$ for $x \in [0,1]$.

PROOF. It suffices to show the first claim, which follows by making a change of variables $t = x^{2k}$ so that we get

$$\int_{-1}^{1} g(x^{2k})dx = 2\int_{0}^{1} g(x^{2k})dx = 2\int_{0}^{1} g(t)\frac{t^{-1+\frac{1}{2k}}}{2k}dt = \frac{1}{k}\int_{0}^{1} g(t)t^{-1+\frac{1}{2k}}dt.$$

□

PROOF OF THEOREM 4.10. By applying the change of variables $x = 2t-1$, we see that the Jacobi type measure $(1-x)^a(1+x)^b dx$ on $[-1,1]$ corresponds to the measure $2^{a+b}(1-t)^a t^b dt$ on $[0,1]$ and that (up to scaling) the orthogonal polynomials for the latter measure on $[0,1]$ are given by $t \mapsto p_i^{a,b}(2t-1)$ for $i \in \mathbb{N}$.

If we set $a = 0$ and $b = -1 + 1/2k$, then the measure obtained in this way on $[0,1]$ is precisely the push-forward measure μ_k (see Lemma 4.12). Hence, we can conclude that (up to scaling) the orthogonal polynomials $p_{k,i}$ for μ_k on $[0,1]$ are given by $p_{k,i}(t) = \mathcal{J}_i^{(a,b)}(2t-1)$ for each $i \in \mathbb{N}$. Therefore, the smallest root of $p_{k,r+1}(t)$ is equal to $(\xi_{r+1}^{a,b}+1)/2 = \Theta(1/r^2)$ by (4.15). In particular, we can conclude that $\text{ub}(f,[-1,1])_r^{\text{pf}} = \Omega(1/r^2)$ for any $k \geq 1$. □

4.4. On the geometric assumption

As mentioned above, the condition (4.4) is a weaker version of a condition introduced in [**dKLS17**]. There, the authors demand that there exist constants $\eta_{\mathbf{X}}, \varepsilon_{\mathbf{X}}$ such that

$$\text{vol}(\mathbf{X} \cap B_\delta^n(\mathbf{x})) \geq \eta_{\mathbf{X}}\delta^n \,\text{vol}(B^n) \quad \forall\, \mathbf{x} \in \mathbf{X},\ \forall\, 0 < \delta \leq \varepsilon_{\mathbf{X}}. \tag{4.16}$$

The difference is that the power of δ in (4.16) is fixed to be the dimension n of $\mathbf{X}$, whereas it is allowed to be an arbitrary $N \geq n$ in (4.4).

Condition (4.4) is satisfied by a significantly larger class of sets $\mathbf{X}$ than (4.16). In particular, as we will observe below, sets satisfying (4.4) may have *polynomial* cusps, whereas sets satisfying (4.16) may not have any cusps at all.

EXAMPLE 4.13. *Consider the set $\mathbf{X} = \{\mathbf{x} \in \mathbb{R}^2 : 0 \leq \mathbf{x}_1 \leq 1,\ 0 \leq \mathbf{x}_2 \leq \mathbf{x}_1^2\}$ (see Figure 4.3). This set $\mathbf{X}$ satisfies* (4.4) *(with $N=3$), but it does not satisfy* (4.16)*. Indeed, for the point $0 \in \mathbf{X}$ we have*

$$\text{vol}(\mathbf{X} \cap B_\delta^2(0)) \leq \int_0^\delta t^2 dt = \frac{1}{3}\delta^3,$$

and conclude (4.16) *cannot be satisfied at* $\mathbf{x} = 0$. *Note that the point* 0 *is indeed a polynomial cusp of the set* $\mathbf{X}$.

EXAMPLE 4.14. *Consider the set* $\mathbf{X} = \{\mathbf{x} \in \mathbb{R}^2 : 0 \leq \mathbf{x}_1 \leq 1,\ 0 \leq \mathbf{x}_2 \leq \exp(-1/\mathbf{x}_1)\}$ *(see Figure 4.3). This set* $\mathbf{X}$ *does not satisfy* (4.4) *(and, as a consequence, does not satisfy* (4.16)*). Indeed, for the point* $0 \in \mathbf{X}$ *we have*

$$\mathrm{vol}(\mathbf{X} \cap B^2_\delta(0)) \leq \int_0^\delta \exp(-1/t)dt \leq \delta \exp(-1/\delta).$$

Now note that for any $N, \eta > 0$ *fixed, we have:*

$$\lim_{\delta \to 0} \frac{\eta \delta^N}{\delta \exp(-1/\delta)} = \infty,$$

and so (4.4) *can not be satisfied at* $\mathbf{x} = 0$. *Note that the point* 0 *is an exponential cusp of* $\mathbf{X}$.

It turns out that compact semialgebraic sets which have a dense interior (aka being fat) satisfy Assumption 4.1, as is shown essentially in [**WP86**].

DEFINITION 4.15. *A set* $\mathbf{X} \subseteq \mathbb{R}^n$ *is called* fat *if* $\mathbf{X} \subseteq \overline{\mathrm{int}(\mathbf{X})}$, *i.e., the interior of* $\mathbf{X}$ *is dense in* $\mathbf{X}$.

THEOREM 4.16 ([**WP86**], Theorem 6.4, see also Remark 6.5). *Let* $\mathbf{X} \subseteq \mathbb{R}^n$ *be a compact, fat semialgebraic*[1] *set. Then there exist constants* $\eta > 0$, $N \geq 1$ *and a positive integer* $d \in \mathbb{N}$ *such that one may find a polynomial* $h_{\mathbf{x}}$ *of degree* d *for each* $\mathbf{x} \in \mathbf{X}$ *satisfying:*

$$\begin{aligned} &h_{\mathbf{x}}(0) = \mathbf{x}, \\ &h_{\mathbf{x}}(t) \in \mathbf{X} \text{ for } t \in [0,1], \text{ and} \\ &B^n_{\eta t^N}(h_{\mathbf{x}}(t)) \subseteq \mathbf{X} \text{ for } t \in [0,1]. \end{aligned} \tag{4.17}$$

Furthermore, the polynomials $h_{\mathbf{x}}$ *may be chosen such that* $\|\mathbf{x} - h_{\mathbf{x}}(t)\| \leq t$ *for all* $\mathbf{x} \in \mathbf{X}, t \in [0,1]$.

COROLLARY 4.17. *Let* $\mathbf{X} \subseteq \mathbb{R}^n$ *be a compact, fat semialgebraic set. Then* $\mathbf{X}$ *satisfies Assumption 4.1.*

PROOF. For $\mathbf{x} \in \mathbf{X}$, let η, N and $h_{\mathbf{x}}$ be as in Theorem 4.16. We may assume that $h_t := h_{\mathbf{x}}(t) \in B^n_t(\mathbf{x})$ for all $t \in [0,1]$. For clarity, we write $B(\mathbf{y}, a) := B^n_a(\mathbf{y})$ in the rest of the proof.

Using the triangle inequality and (4.17) we find that

$$\mathrm{vol}\, B(h_t, \eta t^N) \leq \mathrm{vol}\left(B(\mathbf{x}, t + \eta t^N) \cap \mathbf{X}\right) \leq \mathrm{vol}\left(B(\mathbf{x}, (1+\eta)t) \cap \mathbf{X}\right)$$

[1]In fact, the result is shown for *subanalytic sets*, of which semialgebraic sets are an example.

for all $t \in [0,1]$, noting that $t^N \leq t$ in this case. But now substituting $\delta = (1+\eta)t$ yields:

$$\mathrm{vol}\,\big(B(\mathbf{x},\delta)\cap \mathbf{X}\big) \geq \mathrm{vol}\,B(h_t, \eta\delta^N(1+\eta)^{-N}) = \left(\frac{\eta}{(1+\eta)^N}\right)^n \cdot \delta^{Nn}\,\mathrm{vol}\,B(\mathbf{x},1),$$

showing (4.4) holds for $0 < \delta \leq \varepsilon_{\mathbf{X}} := (1+\eta)$ and $\eta_{\mathbf{X}} = \eta^n(1+\eta)^{-Nn}$. $\square$

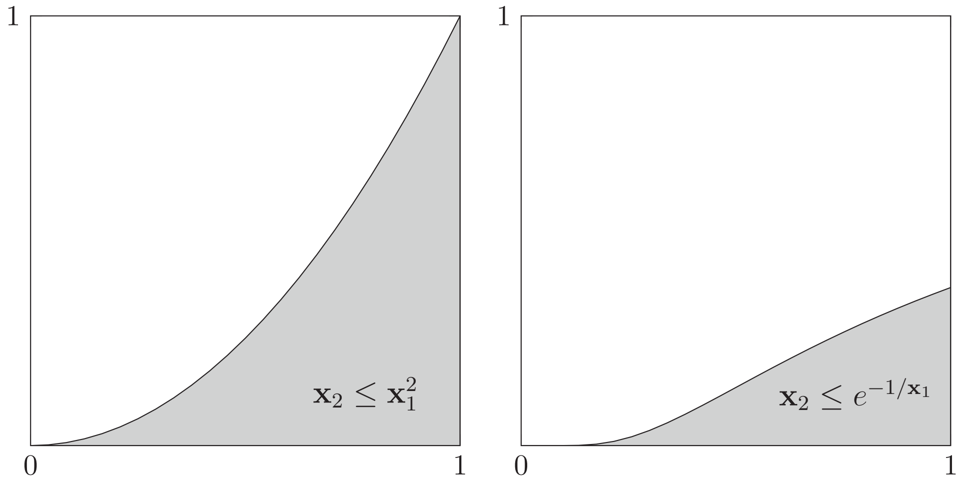

FIGURE 4.3. The set $\mathbf{X} = \{\mathbf{x} \in \mathbb{R}^2 : 0 \leq \mathbf{x}_1 \leq 1, 0 \leq \mathbf{x}_2 \leq \mathbf{x}_1^2\}$ (left) and $\mathbf{X} = \{\mathbf{x} \in \mathbb{R}^2 : 0 \leq \mathbf{x}_1 \leq 1, 0 \leq \mathbf{x}_2 \leq \exp(-1/\mathbf{x}_1)\}$ (right).

4.5. Discussion

We have shown a convergence rate in $O(\log^2 r/r^2)$ for the approximations $\mathrm{ub}(f,\mathbf{X})_r^{\mathrm{pf}}$ of the minimum of a polynomial f over a compact set $\mathbf{X}$ satisfying the minor geometric assumption (4.4). This includes in particular fat semi-algebraic sets. As a direct consequence this rate also holds for the regular measure-based bounds $\mathrm{ub}(f,\mathbf{X})_r$. Furthermore, we have shown that this analysis is near-optimal, in the sense that the asymptotic behaviour of the error range $\mathrm{ub}(f,\mathbf{X})_r^{\mathrm{pf}} - f_{\min}$ is in $O(\log^2 r/r^2)$ in general and in $\Omega(1/r^2)$ for an infinite class of polynomials.

In fact, a rate in $O(\log^2 r/r^2)$ for the bounds $\mathrm{ub}(f,\mathbf{X})_r$ is established *directly* in our work [**SL20**]. The argument there makes use of both the needle and $\frac{1}{2}$-needle polynomials. It is significantly more complicated than the proof of Theorem 4.9 above, as it relies more intricately on the geometric properties of the set $\mathbf{X}$. Furthermore, the result applies less generally than Theorem 4.9;

namely only to convex bodies. For these reasons, we have not included it in this thesis.

One advantage of the (proof of) the result in [**SL20**], however, is that it allows quite easily to show an improved convergence rate in $O(\log^k r/r^k)$ of the bounds $\mathrm{ub}(f, \mathbf{X})_r$ when the first k derivatives of f vanish at its global minimizer. Using this fact, we showed in Section 4.3 that although the worst-case guarantees on the convergence of the regular measure-based bounds $\mathrm{ub}(f, \mathbf{X})_r$ and push-forward bounds $\mathrm{ub}(f, \mathbf{X})_r^{\mathrm{pf}}$ are very similar, a large separation may exist for certain polynomials (e.g., when $f(x) = x^{2k}$, $k \geq 2$).

The main question left open is whether the log-factor in our main result can be avoided. This factor arises from our analysis technique, and in particular from our use of the needle polynomials to approximate the Dirac function on $[0, 1]$. Recall from Chapter 3 that a rate in $O(1/r^2)$ may be established for the bounds $\mathrm{ub}(f, \mathbf{X}, \mu)_r$ in several (fundamental) special cases. To obtain such rates, one ultimately relies on understanding the behaviour of certain orthogonal polynomials on the interval $[-1, 1]$ w.r.t. the Chebyshev measure (cf. Chapter 2). In the push-forward setting, the relevant orthogonal polynomials are much more poorly understood. Therefore, it seems hard to establish a rate in $O(1/r^2)$ for the bounds $\mathrm{ub}(f, \mathbf{X}, \mu)_r^{\mathrm{pf}}$, even when $\mathbf{X}$ is a well-structured set.

A second question is whether the upper bounds converge significantly more slowly on sets $\mathbf{X}$ which do not satisfy the geometric assumption (4.4). We suspect that this might be the case, particularly when $\mathbf{X}$ has exponential cusps (see Figure 4.3).

Acknowledgments. We wish to thank Edouard Pauwels for bringing the needle polynomials to our attention.

CHAPTER 5

Computational aspects of measure-based bounds

> *Whenever a man can get hold of numbers, they are invaluable: if correct, they assist in informing his own mind, but they are still more useful in deluding the minds of others.*
>
> Charles Babbage

This chapter is partly based on my joint works [**SL20, SL21a**] *with Monique Laurent.*

In this chapter, we discuss some computational aspects of the measure-based upper bounds $\mathrm{ub}(f)_r$ and $\mathrm{ub}(f)_r^{\mathrm{pf}}$ for the minimization of a polynomial f over a compact (semialgebraic) set $\mathbf{X}$. In Section 5.1, we outline potential challenges in computing the values of these bounds, and offer some suggestions to address them. In Section 5.2, we discuss the problem of obtaining approximate *minimizers* $\widetilde{\mathbf{x}} \approx \mathbf{x}^* \in \mathbf{X}$ of f based on feasible solutions (i.e., sum-of-squares densities) to the programs defining $\mathrm{ub}(f)_r$ and $\mathrm{ub}(f)_r^{\mathrm{pf}}$. Finally, in Section 5.3, we present some numerical examples that illustrate the practical behaviour of these bounds and some of the results of Chapters 3 and 4.

We only introduce a few (minor) new ideas in these sections. Their main purpose is rather to highlight some aspects of the measure-based bounds which are not very well explored in the literature, and offer potentially interesting new directions for future research.

5.1. Computing the measure-based bounds

Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact set and let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree d. We consider the optimization problem:

$$f_{\min} := \min_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x}). \tag{5.1}$$

Let μ be a positive Borel measure supported on $\mathbf{X}$. Recall the definition of Lasserre's measure-based hierarchies of upper bounds on $f_{\min}$:

$$\text{ub}(f, \mathbf{X}, \mu)_r = \inf_{q \in \Sigma[\mathbf{x}]_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mathbf{x} : \int_{\mathbf{X}} q(\mathbf{x}) d\mathbf{x} = 1 \right\}, \tag{5.2}$$

$$\text{ub}(f, \mathbf{X}, \mu)_r^{\text{pf}} = \inf_{u \in \Sigma[x]_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) u(f(\mathbf{x})) d\mathbf{x} : \int_{\mathbf{X}} u(f(\mathbf{x})) d\mathbf{x} = 1 \right\}. \tag{5.3}$$

For convenience, we will often omit the dependence of these bounds on $\mathbf{X}$ and μ in the notation; writing simply $\text{ub}(f)_r$ and $\text{ub}(f)_r^{\text{pf}}$. A feasible solution $u \in \Sigma[x]_{2r}$ to (5.3) corresponds to a feasible solution $q(\mathbf{x}) = u(f(\mathbf{x}))$ to (5.2). That is, for any $r \in \mathbb{N}$, we have:

$$\text{ub}(f)_r^{\text{pf}} \leq \text{ub}(f)_{rd}. \tag{5.4}$$

As we have seen in Section 2.2, the programs (5.2) and (5.3) may both be reduced to an eigenvalue optimization problem. Namely, we have:

$$\text{ub}(f)_r = \lambda_{\min}\big(M_{f,r}(\mu), M_{1,r}(\mu)\big), \tag{5.5}$$

$$\text{ub}(f)_r^{\text{pf}} = \lambda_{\min}\big(M_{x,r}(\mu_f), M_{1,r}(\mu_f)\big). \tag{5.6}$$

Here, we make use of the (truncated) *moment matrices* for the measure μ and for the push-forward measure $\mu_f(\mathbf{x})$ given by $\mu_f(A) = \mu(f^{-1}(A))$. For a polynomial $g \in \mathbb{R}[\mathbf{x}]$ and an integer k, they are defined as:

$$\big(M_{g,r}(\mu)\big)_{\alpha,\beta} = \int_{\mathbf{X}} g(\mathbf{x}) \mathbf{x}^{\alpha} \mathbf{x}^{\beta} d\mu(\mathbf{x}) \qquad (|\alpha|, |\beta| \leq r), \tag{5.7}$$

$$\big(M_{x^k,r}(\mu_f)\big)_{i,j} = \int_{f(\mathbf{X})} x^{i+j+k} d\mu_f(x) = \int_{\mathbf{X}} f(\mathbf{x})^{i+j+k} d\mu(\mathbf{x}) \quad (i, j \leq r). \tag{5.8}$$

The formulations (5.5) and (5.6) (together with the equations (5.7) and (5.8)) reveal the main advantage of the push-forward bounds $\text{ub}(f)_r^{\text{pf}}$: Their computation involves matrices of size $r + 1$, thus much smaller than the matrices of size $\binom{n+r}{r}$ involved in the computation of the regular measure-based bounds $\text{ub}(f)_r$.

5.1.1. Computing moments. A major downside of the upper bounds $\text{ub}(f)_r$ and $\text{ub}(f)_r^{\text{pf}}$ over the more well-known *lower* bounds introduced in Chapter 2 is that their computation requires *explicit* knowledge of the moment matrices (5.7) and (5.8), respectively. From a practical point of view, this excludes many sets $\mathbf{X}$. Indeed, it is already hard in general to estimate the volume of a semialgebraic set [**Las09a**].

In certain special cases, on the other hand, closed form expressions of the moments $\int_{\mathbf{X}} \mathbf{x}^{\alpha} d\mu(\mathbf{x})$ are available. This includes the binary cube, the hypersphere, the unit ball, the box $[-1, 1]^n$ and the standard simplex, see [**dKL19**]. Such expressions allow one to compute the entries of the moment matrices (5.7), (5.8) after writing $f(\mathbf{x}) = \sum_{|\alpha| \leq d} f_{\alpha} \mathbf{x}^{\alpha}$ in the monomial basis and taking linear combinations.

For the regular bounds $\mathrm{ub}(f)_r$ this is straightforward, yet still computationally intensive as the matrices in (5.5) are of large size $\binom{n+r}{r}$. For the push-forward bounds $\mathrm{ub}(f)_r^{\mathrm{pf}}$, only $O(r)$ different moments have to be computed. However, they are now of the form:

$$\int_{\mathbf{X}} f(\mathbf{x})^k d\mu(\mathbf{x}) \quad (k \leq 2r). \tag{5.9}$$

This poses a problem, as computing the k-th power $f(\mathbf{x})^k$ of f (and expressing it in the monomial basis) is not a trivial task. Indeed, the polynomial $f(\mathbf{x})^k$ may have exponentially many terms. Naive computation of (5.9) therefore comes at the risk of losing any computational advantage over the regular (asymptotically stronger) bounds $\mathrm{ub}(f)_r$. This motivates the following question.

Question 5.1. *Suppose we have access to the moments* $\int_{\mathbf{X}} \mathbf{x}^\alpha d\mu(\mathbf{x})$ *for certain* $\mathbf{X} \subseteq \mathbb{R}^n$ *and* μ. *Given a polynomial* $f \in \mathbb{R}[\mathbf{x}]$ *and a number* $k \in \mathbb{N}$, *what is the most efficient way to (approximately) compute the integral:*

$$\int_{\mathbf{X}} f(\mathbf{x})^k d\mu(\mathbf{x})?$$

5.1.2. Symmetry reduction. The role of symmetry and sparsity in the computation of the *lower* bounds $\mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r$ (and its variants) is well explored in the literature. For the upper bounds, on the other hand, little is known.

We make here a small observation which could potentially be exploited in this setting. Suppose $G \subseteq O(n)$ is a compact group acting on $\mathbb{R}^n$ in the usual way. We say that $\mathbf{X} \subseteq \mathbb{R}^n$ is invariant under G if $\pi(\mathbf{X}) = \mathbf{X}$ for all $\pi \in G$. Similarly, we say a Borel measure μ on G is invariant if $\mu(\pi(A)) = \mu(A)$ for all $A \subseteq \mathbf{X}$ measurable. Finally, a polynomial $f \in \mathbb{R}[\mathbf{x}]$ is invariant if $f(\mathbf{x}) = f(\pi(\mathbf{x}))$ for all $\mathbf{x} \in \mathbb{R}^n$ and $\pi \in G$.

Proposition 5.2. *Let* $G \subseteq O(n)$ *compact. Assume that* $\mathbf{X}, \mu$ *and* f *are all* G*-invariant. Then there exists an optimum solution* $q \in \Sigma[\mathbf{x}]_{2r}$ *to the formulation* (5.2) *of* $\mathrm{ub}(f, \mathbf{X}, \mu)_r$ *which is* G*-invariant.*

Proof. Let $\widetilde{q}$ be *any* optimum solution to (5.2). Let ω be the uniform probability measure on G, and consider the G-invariant density $q \in \Sigma[\mathbf{x}]_{2r}$ given by:

$$q(\mathbf{x}) = \int_G \widetilde{q}(\pi(\mathbf{x})) d\omega(\pi).$$

Using the invariance of $\mathbf{X}, \mu$ and f, we find that:

$$\int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mu(\mathbf{x}) = \int_G \int_{\mathbf{X}} f(\mathbf{x}) \widetilde{q}(\pi(\mathbf{x})) d\mu(\mathbf{x}) d\omega(\pi) = \int_{\mathbf{X}} f(\mathbf{x}) \widetilde{q}(\mathbf{x}) d\mu(\mathbf{x}).$$

Therefore, q is also an optimum solution to (5.2), as desired. □

5.1.3. Parallelization. Finally, we wish to mention that the computation of the upper bounds is particularly suited to *parallelization*. Solving the eigenvalue formulations (5.5) and (5.6) requires a) computing the entries of the appropriate moment matrices and b) solving the resulting (generalized) eigenvalue problem. Both tasks can be performed in parallel after a suitable subdivision of the rows and/or columns of the involved matrices, assigning a block to each computational node. Indeed, one may iteratively compute eigenvalues using only matrix-vector multiplications. These matrix-vector multiplications are straightforwardly parallelizable.

In some very preliminary experiments, this approach leads to reasonable speedups for the bounds $\mathrm{ub}(f)_r$. For the push-forward bounds $\mathrm{ub}(f)_r^{\mathrm{pf}}$, a naive parallel implementation does not seem to offer much improvement. As relatively few different moments $\int_{\mathbf{X}} f(\mathbf{x})^k d\mu(\mathbf{x})$ are involved there, one should probably focus instead on computing those integrals in parallel, see also Question 5.1 above.

5.2. Extracting minimizers

In this section we discuss a method for extracting a feasible solution $\widetilde{\mathbf{x}} \in \mathbf{X}$ to the minimization problem (5.1) from an optimum density $q_r \in \Sigma[\mathbf{x}]_{2r}$ to the relaxation (5.2). The idea behind this method is already mentioned (in a slightly different context) in the work [**dKLLS17**]. In principle, we would like such solutions to satisfy the following properties:

(1) $f(\widetilde{\mathbf{x}}) \approx \mathrm{ub}(f)_r$;
(2) $\widetilde{\mathbf{x}} \approx \mathbf{x}^*$, meaning that $\|\widetilde{\mathbf{x}} - \mathbf{x}^*\|$ is small for some minimizer $\mathbf{x}^* \in \mathbf{X}$ of f;
(3) $\widetilde{\mathbf{x}} \in \mathbf{X}$, meaning that $\widetilde{\mathbf{x}}$ is a feasible solution to (5.1).

For simplicity, we assume throughout that we are in a setting where $\lim_{r\to\infty} \mathrm{ub}(f)_r = f_{\min}$, i.e., where the measure-based bounds converge to the optimum value of (5.1). We note that the optimum solution to (5.2) (and thus the density q_r) need not be unique. One should therefore interpret the results and statements below as being true *independently* of the choice of optimum solution.

5.2.1. A feasible point. For $r \in \mathbb{N}$, consider the point $\widetilde{\mathbf{x}} = \mathbf{x}^{(r)} \in \mathbb{R}^n$, defined in terms of an optimum density $q_r \in \Sigma[\mathbf{x}]_{2r}$ as:

$$\mathbf{x}_i^{(r)} := \int_{\mathbf{X}} \mathbf{x}_i q_r(\mathbf{x}) d\mu(\mathbf{x}) \quad (1 \leq i \leq n).$$

An immediate advantage of this approach is that $\mathbf{x}^{(r)}$ may easily be computed as long as the moments of μ are known. A second advantage is that $\mathbf{x}^{(r)}$ is a feasible solution to (5.1) as long as $\mathbf{X}$ is convex.

PROPOSITION 5.3 ([**dKLLS17**]). *If $\mathbf{X}$ is convex, then $\mathbf{x}^{(r)} \in \mathbf{X}$ for all $r \in \mathbb{N}$.*

Finally, if the polynomial f is also assumed to be convex on $\mathbf{X}$, the value $f(\mathbf{x}^{(r)})$ of f at $\mathbf{x}^{(r)}$ is at most equal to the upper bound $\mathrm{ub}(f)_r$. This follows readily after applying Jensen's inequality.

Proposition 5.4 ([**dKLLS17**]). *Assume that* $\mathbf{X}$ *is convex and that* f *is convex on* $\mathbf{X}$. *Then for each* $r \in \mathbb{N}$, *we have:*

$$f(\mathbf{x}^{(r)}) \leq \mathrm{ub}(f)_r.$$

A consequence of Proposition 5.4 is that $\mathbf{x}^{(r)}$ converges to the set of minimizers of f.

Corollary 5.5. *Assume that* $\mathbf{X} \subseteq \mathbb{R}^n$ *is convex and that* f *is convex on* $\mathbf{X}$. *Write* $\mathbf{X}^* = \{\mathbf{x} \in \mathbf{X} : f(\mathbf{x}) = f_{\min}\}$ *for the set of minimizers of* f *on* $\mathbf{X}$. *Then we have:*

$$\lim_{r \to \infty} d_\infty(\mathbf{x}^{(r)}, \mathbf{X}^*) = 0.$$

Here, $d_\infty(\mathbf{x}^{(r)}, \mathbf{X}^*) := \min_{\mathbf{x}^* \in \mathbf{X}^*} \|\mathbf{x}^{(r)} - \mathbf{x}^*\|_\infty$.

Proof. As both $\mathbf{X}$ and f are convex, the set $\mathbf{X}^*$ is convex as well. Furthermore, $\mathbf{X}^* \subseteq \mathbf{X}$ is closed and thus compact (as $\mathbf{X}$ is compact). Now suppose that $\lim_{r\to\infty} d_\infty(\mathbf{x}^{(r)}, \mathbf{X}^*) \neq 0$. Then there exists an $\eta > 0$ and a subsequence $(r_i)_{i\in\mathbb{N}}$ such that $d_\infty(\mathbf{x}^{(r_i)}, \mathbf{X}^*) \geq \eta$ for all $i \in \mathbb{N}$. Let $\mathbf{X}^*_\eta := \{\mathbf{x} \in \mathbf{X} : d_\infty(\mathbf{x}, \mathbf{X}^*) \geq \eta\}$. Then $\mathbf{X}^*_\eta$ is a compact subset of $\mathbf{X}$ which does not intersect $\mathbf{X}^*$, and so:

$$\min_{\mathbf{x} \in \mathbf{X}^*_\eta} f(\mathbf{x}) > f_{\min}.$$

We may conclude that $\liminf_{r\to\infty} f(\mathbf{x}^{(r)}) > f_{\min}$, which contradicts Proposition 5.4 (recalling that we assume that $\lim_{r\to\infty} \mathrm{ub}(f)_r = f_{\min}$). □

The assumption that f is convex on $\mathbf{X}$ is rather prohibitive. In the rest of this section, we therefore explore what happens when we drop this assumption.

5.2.2. A negative result. We begin with a negative result showing that, in general, we cannot hope that $\mathbf{x}^{(r)}$ approximates a minimizer $\mathbf{x}^*$ of f if we drop the convexity assumption on f. For this purpose, consider the interval $\mathbf{X} = [-1, 1]$, equipped with the Lebesgue measure and set $f(x) = 1 - x^2$. Then the minimum $f_{\min} = 0$ of f on $\mathbf{X}$ is attained at $x^* \in \{-1, 1\}$.

Lemma 5.6. *Let* $r \in \mathbb{N}$. *There exists an optimum solution* $q_r \in \Sigma[x]_{2r}$ *to the program* (5.2) *which is an* even *function, i.e., satisfying* $q_r(x) = q_r(-x)$ *for* $x \in \mathbb{R}$.

Proof. Let $\widetilde{q}_r \in \Sigma[x]_{2r}$ be *any* optimum solution to (5.2). Consider the polynomial $q_r \in \mathbb{R}[x]$ given by:

$$q_r(x) = \frac{\widetilde{q}_r(x) + \widetilde{q}_r(-x)}{2}.$$

Clearly, $q_r \in \Sigma[x]_{2r}$ is an even function. Furthermore, using the fact that $f(x) = 1 - x^2$ is even, we have:

$$\int_{-1}^{1} f(x)q_r(x)dx = \int_{-1}^{1} f(x)\widetilde{q}_r(x)dx = \text{ub}(f)_r.$$

Thus, q_r is an optimum solution to (5.1), as desired. □

Now suppose that we select our optimum density q_r for each r as in Lemma 5.6. Then, since q_r is even, the candidate solution $\mathbf{x}^{(r)}$ is given for each r by:

$$\mathbf{x}^{(r)} = \int_{-1}^{1} xq_r(x)dx = 0.$$

We thus obtain no information on the minimizers $\{-1, 1\}$ of f in this case.

5.2.3. A positive result. As the following proposition shows, when f has a unique minimizer $\mathbf{x}^*$ on $\mathbf{X}$, asymptotic convergence of $\mathbf{x}^{(r)}$ to $\mathbf{x}^*$ is guaranteed. This may be seen as a generalization of Corollary 5.5.

PROPOSITION 5.7. *If $\mathbf{X}$ is compact and f has a unique minimizer $\mathbf{x}^* \in \mathbf{X}$, then we have:*

$$\lim_{r\to\infty} \mathbf{x}^{(r)} = \mathbf{x}^*.$$

We recall that $\mathbf{x}^{(r)}$ is defined in terms of an optimum density $q_r \in \Sigma[\mathbf{x}]_{2r}$ in (5.2) by $\mathbf{x}_i^{(r)} = \int_{\mathbf{X}} \mathbf{x}_i q_r(\mathbf{x})d\mu(\mathbf{x})$, $1 \leq i \leq n$.

LEMMA 5.8. *For $\mathbf{x} \in \mathbb{R}^n$ and $\eta \in \mathbb{R}$, let $D(\mathbf{x}, \eta) := \{\mathbf{y} \in \mathbb{R}^n : \|\mathbf{x} - \mathbf{y}\|_\infty < \eta\}$. Let $\mathbf{X} \subseteq \mathbb{R}^n$ be compact and assume that f has unique minimizer $\mathbf{x}^*$ on $\mathbf{X}$. Then for any $\eta > 0$, the optimum densities q_r in (5.2) satisfy:*

$$\lim_{r\to\infty} \int_{\mathbf{X}\setminus D(\mathbf{x}^*,\eta)} q_r(\mathbf{x})d\mu(\mathbf{x}) = 0.$$

PROOF. Suppose not. Then there exists a subsequence $(r_i)_{i\in\mathbb{N}}$ and an $\varepsilon > 0$ so that:

$$\int_{\mathbf{X}\setminus D(\mathbf{x}^*,\eta)} q_{r_i}(\mathbf{x})d\mu(\mathbf{x}) > \varepsilon$$

for all $i \in \mathbb{N}$. By compactness of $\mathbf{X}$ and uniqueness of the minimizer $\mathbf{x}^*$, we have:

$$\min_{\mathbf{x}\in\mathbf{X}\setminus D(\mathbf{x}^*,\eta)} f(\mathbf{x}) > f_{\min}.$$

But this implies that:

$$\begin{aligned}
\mathrm{ub}(f)_{r_i} &= \int_{\mathbf{X}} q_{r_i}(\mathbf{x}) f(\mathbf{x}) d\mu(\mathbf{x}) \\
&= \int_{D(\mathbf{x}^*,\eta)} q_{r_i}(\mathbf{x}) f(\mathbf{x}) + \int_{\mathbf{X}\setminus D(\mathbf{x}^*,\eta)} q_{r_i}(\mathbf{x}) f(\mathbf{x}) \\
&> (1-\varepsilon) f_{\min} + \varepsilon \cdot \min_{\mathbf{x}\in\mathbf{X}\setminus D(\mathbf{x}^*,\eta)} f(\mathbf{x}), \\
&= f_{\min} + \varepsilon \cdot \big(\min_{\mathbf{x}\in\mathbf{X}\setminus D(\mathbf{x}^*,\eta)} f(\mathbf{x}) - f_{\min} \big)
\end{aligned}$$

for all $i \in \mathbb{N}$, contradicting the fact that $\lim_{r\to\infty} \mathrm{ub}(f)_r = f_{\min}$. □

Proof of Proposition 5.7. Let $\varepsilon > 0$ and fix an $i \in [n]$. As $\mathbf{X}$ is compact, there exists $\rho > 0$ such that $|\mathbf{x}_i - \mathbf{x}_i^*| < \rho$ for all $\mathbf{x} \in \mathbf{X}$. For $r \in \mathbb{N}$, let $q_r \in \Sigma[\mathbf{x}]_{2r}$ be an optimum density in (5.2). Then we have:

$$\begin{aligned}
\lim_{r\to\infty} |\mathbf{x}_i^{(r)} - \mathbf{x}_i^*| &= \lim_{r\to\infty} \Big| \int_{\mathbf{X}} (\mathbf{x}_i - \mathbf{x}_i^*) q_r(\mathbf{x}) d\mu(\mathbf{x}) \Big| \\
&\leq \lim_{r\to\infty} \Big| \int_{D(\mathbf{x}^*,\varepsilon)} (\mathbf{x}_i - \mathbf{x}_i^*) q_r(\mathbf{x}) d\mu(\mathbf{x}) \Big| \\
&\quad + \Big| \int_{\mathbf{X}\setminus D(\mathbf{x}^*,\varepsilon)} (\mathbf{x}_i - \mathbf{x}_i^*) q_r(\mathbf{x}) d\mu(\mathbf{x}) \Big| \\
&\leq \varepsilon + \rho \lim_{r\to\infty} \Big| \int_{\mathbf{X}\setminus D(\mathbf{x}^*,\varepsilon)} (\mathbf{x}_i - \mathbf{x}_i^*) q_r(\mathbf{x}) d\mu(\mathbf{x}) \Big| = \varepsilon,
\end{aligned}$$

making use of Lemma 5.8 (with $\eta = \varepsilon$) for the last equality. We have thus shown that $\lim_{r\to\infty} \mathbf{x}^{(r)} = \mathbf{x}^*$. □

5.2.4. A quantitative result. We may bound the rate of convergence of $\mathbf{x}^{(r)}$ to $\mathbf{x}^*$ in terms of the convergence rate of $\mathrm{ub}(f)_r$ to $f_{\min}$. This rate will depend on the local behaviour of f near $\mathbf{x}^*$.

Proposition 5.9. *Assume that $\mathbf{X}$ is compact and that f has a unique minimizer $\mathbf{x}^* \in \mathbf{X}$. Assume further that there exist $c > 0$ and $k \in \mathbb{N}$ such that:*

$$f(\mathbf{x}) \geq c\|\mathbf{x} - \mathbf{x}^*\|^k + f(\mathbf{x}^*) \quad (\mathbf{x} \in \mathbf{X}).$$

Then we have:

$$\|\mathbf{x}^{(r)} - \mathbf{x}^*\|_\infty \leq 2 \left(\frac{\mathrm{diam}_\infty(\mathbf{X})}{c} \cdot (\mathrm{ub}(f)_r - f_{\min}) \right)^{\frac{1}{k+1}}.$$

Here, $\mathrm{diam}_\infty(\mathbf{X}) := \max_{\mathbf{x},\mathbf{y}\in\mathbf{X}} \|\mathbf{x} - \mathbf{y}\|_\infty < \infty$.

Lemma 5.10 (See also Lemma 3.2 in [**SL20**]). *Suppose the conditions of Proposition 5.9 are met. Then, for any $\varepsilon > 0$, we have:*

$$\int_{\mathbf{X}\setminus D(\mathbf{x}^*,\varepsilon)} q_r(\mathbf{x}) d\mu(\mathbf{x}) \leq \frac{\mathrm{ub}(f)_r - f_{\min}}{c\varepsilon^k}.$$

Proof. By assumption, we have $f(\mathbf{x}) \geq c\varepsilon^k + f_{\min}$ for all $\mathbf{x} \notin D(\mathbf{x}^*, \varepsilon)$. Therefore, we find that:

$$\mathrm{ub}(f)_r - f_{\min} = \int_{\mathbf{X}} (f(\mathbf{x}) - f_{\min}) q_r(\mathbf{x}) d\mu(\mathbf{x}) \geq c\varepsilon^k \int_{\mathbf{X} \setminus D(\mathbf{x}^*, \varepsilon)} q_r(\mathbf{x}) d\mu(\mathbf{x}).$$

□

Proof of Proposition 5.9. Let $\varepsilon > 0$ and fix $i \in [n]$. We have that:

$$\begin{aligned}
|\mathbf{x}_i^{(r)} - \mathbf{x}_i^*| &= |\int_{\mathbf{X}} (\mathbf{x}_i - \mathbf{x}_i^*) q_r(\mathbf{x}) d\mu(\mathbf{x})| \\
&\leq |\int_{D(\mathbf{x}^*, \varepsilon)} (\mathbf{x}_i - \mathbf{x}_i^*) q_r(\mathbf{x}) d\mu(\mathbf{x})| + |\int_{\mathbf{X} \setminus D(\mathbf{x}^*, \varepsilon)} (\mathbf{x}_i - \mathbf{x}_i^*) q_r(\mathbf{x}) d\mu(\mathbf{x})| \\
&\leq \varepsilon + \mathrm{diam}_\infty(\mathbf{X}) \int_{\mathbf{X} \setminus D(\mathbf{x}^*, \varepsilon)} q_r(\mathbf{x}) d\mu(\mathbf{x}) \\
&\leq \varepsilon + \mathrm{diam}_\infty(\mathbf{X}) \frac{\mathrm{ub}(f)_r - f_{\min}}{c\varepsilon^k},
\end{aligned}$$

making use of Lemma 5.10 for the last inequality. Now set:

$$\varepsilon = \Big(\frac{\mathrm{diam}_\infty(\mathbf{X})}{c} (\mathrm{ub}(f)_r - f_{\min})\Big)^{1/(k+1)}$$

to finish the proof. □

We note that the rate we obtain in Proposition 5.9 for the convergence of $\mathbf{x}^{(r)}$ to $\mathbf{x}^*$ is likely far from best-possible.

5.2.5. Sampling from the optimum density. One may also obtain feasible points $\widetilde{\mathbf{x}} \in \mathbf{X}$ by *sampling* $\widetilde{\mathbf{x}}$ according to an optimum density $q_r \in \Sigma[\mathbf{x}]_{2r}$ coming from the relaxation (5.2). This idea is explored in some detail in [**dKLS17, dKLLS17**]. The main advantage of this approach is that such points $\widetilde{\mathbf{x}}$ satisfy $\mathbb{E}[\widetilde{\mathbf{x}}] \leq \mathrm{ub}(f)_r$ (without the need of any convexity assumptions). Indeed, this is by definition of the program (5.2). Furthermore, in certain cases, one may sample relatively efficiently using the *method of conditional distributions*, see [**dKLS17**] for details. Therefore, one may obtain a good feasible point $\widetilde{\mathbf{x}}$ with high probability after generating several samples.

5.3. Numerical examples

In this section, we illustrate the practical behaviour of the bounds $\mathrm{ub}(f)_r$ and $\mathrm{ub}(f)_r^{\mathrm{pf}}$ using some numerical examples. Following earlier works [**dKLS17, dKL18, dKLLS17**], we mostly consider the test functions listed in Table 5.1, which are well-known in optimization. We compute the bounds using the eigenvalue reformulations, and solve the resulting eigenvalue problems using the linear algebra module of the SciPy package [**VGO$^+$20**].

Name	Formula	$f_{\min}$
Booth	$f_{\text{bo}}(\mathbf{x}) = (10\mathbf{x}_1 + 20\mathbf{x}_2 - 7)^2 + (20\mathbf{x}_1 + 10\mathbf{x}_2 - 5)^2$	$f_{\text{bo}}(\frac{1}{10}, \frac{3}{10}) = 0$
Matyas	$f_{\text{ma}}(\mathbf{x}) = 26(\mathbf{x}_1^2 + \mathbf{x}_2^2) - 48\mathbf{x}_1\mathbf{x}_2$	$f_{\text{ma}}(0, 0) = 0$
Camel	$f_{\text{ca}}(\mathbf{x}) = 50\mathbf{x}_1^2 - \frac{2625}{4}\mathbf{x}_1^4 + \frac{15625}{6}\mathbf{x}_1^6 + 25\mathbf{x}_1\mathbf{x}_2 + 25\mathbf{x}_2^2$	$f_{\text{ca}}(0, 0) = 0$
Motzkin	$f_{\text{mo}}(\mathbf{x}) = 64\mathbf{x}_1^4\mathbf{x}_2^2 + 64\mathbf{x}_1^2\mathbf{x}_2^4 - 48\mathbf{x}_1^2\mathbf{x}_2^2 + 1$	$f_{\text{mo}}(\pm\frac{1}{2}, \pm\frac{1}{2}) = 0$

TABLE 5.1. Polynomial test functions. In each case, $f_{\min}$ is the global minimum of f on $[-1, 1]^2$.

5.3.1. Comparison of $\text{ub}(f)_r$ and $\text{ub}(f)_r^{\text{pf}}$ for polynomial test functions. First, we compare the bounds $\text{ub}(f)_r$ and $\text{ub}(f)_r^{\text{pf}}$ directly for the test functions $f \in \{f_{\text{bo}}, f_{\text{ma}}, f_{\text{ca}}, f_{\text{mo}}\}$ on the unit box $[-1, 1]^2$ and the unit ball B^2 (w.r.t. the Lebesgue measure). Recall from (5.4) that $\text{ub}(f)_{rd} \leq \text{ub}(f)_r^{\text{pf}}$ for $d = \deg(f)$, i.e., that the regular bound $\text{ub}(f)_r$ is (asymptotically) stronger than the push-forward bound $\text{ub}(f)_r^{\text{pf}}$.

For $1 \leq r \leq 20$, we compute the values of the fraction:

$$\text{ratio}(f)_r := \frac{\text{ub}(f)_r^{\text{pf}} - f_{\min}}{\text{ub}(f)_r - f_{\min}}.$$

So, values of $\text{ratio}(f)_r$ smaller than 1 indicate good performance of the bounds $\text{ub}(f)_r^{\text{pf}}$ in comparison to $\text{ub}(f)_r$. The results can be found in Figure 5.3. Remarkably, it appears that the performance of the bound $\text{ub}(f)_r^{\text{pf}}$ is comparable to (or better than) the performance of $\text{ub}(f)_r$ in each instance, except for the Camel function. Additionally, we note that the performance of $\text{ub}(f)_r^{\text{pf}}$ for the Motzkin polynomial is comparatively much better on the unit ball than on the unit box. Figure 5.1 shows a plot of the Camel function, as well as the sum-of-squares densities corresponding to $\text{ub}(f)_6$ and $\text{ub}(f)_6^{\text{pf}}$ on the unit box.

Note that while the density corresponding to $\text{ub}(f)_6$ resembles the Dirac delta function centered at the global minimizer $(0, 0)$ of the Camel function, the density corresponding to $\text{ub}(f)_6^{\text{pf}}$ instead mirrors the Camel function itself. This is not suprising, as the densities considered in the program (5.3) defining $\text{ub}(f)_r^{\text{pf}}$ are of the form $q(\mathbf{x}) = u(f_{\text{ca}}(\mathbf{x}))$.

5.3.2. Comparison of $\text{ub}(f)_r$ and $\text{ub}(f)_r^{\text{pf}}$ for a special class of polynomials. Next, we consider the polynomials $f(x) = x^{2k}$ for $k \geq 1$ on the interval $[-1, 1]$, which were treated in Section 4.3. Recall that this class of polynomials was used to show a large separation between the bounds $\text{ub}(f)_r$ and $\text{ub}(f)_r^{\text{pf}}$; namely we showed that for these polynomials, we have:

$$\text{ratio}(f)_r := \frac{\text{ub}(f)_r^{\text{pf}} - f_{\min}}{\text{ub}(f)_r - f_{\min}} = \Omega(r^{2k-3}).$$

In Figure 5.4, the values of $\mathrm{ratio}(f)_r$ are shown for $1 \leq r \leq 20$ and $1 \leq k \leq 5$. It can be seen that the performance of $\mathrm{ub}(f)_r^{\mathrm{pf}}$ is comparable to the performance of $\mathrm{ub}(f)_r$ for $k = 1$ (indeed, in this case we have $\mathrm{ub}(f)_r^{\mathrm{pf}} = \mathrm{ub}(f)_{2r}$), but it is much worse for $k > 1$, which matches our earlier findings in Section 4.3. In Figure 5.2, the optimal sum-of-squares densities σ (corresponding to $\mathrm{ub}(f)_r$) and σ_{pf} (corresponding to $\mathrm{ub}(f)_r^{\mathrm{pf}}$) are depicted for $k = 1, 3, 5$ and $r = 6$. Note that while the density σ changes very little as we increase k, the density σ_{pf} grows increasingly 'flat' around the minimizer 0 of f (mirroring the behavior of f itself). As such, the density σ_{pf} is a comparatively much worse approximation of the Dirac delta function centered at 0 than σ. Note also that in this instance $\mathrm{ub}(f)_r = \mathrm{ub}(f)_{r+1}$ for even r, explaining the 'zig-zagging' behaviour of the ratio $\mathrm{ratio}(f)_r$.

Comparison of $\mathrm{ub}(f)_r^{\mathrm{pf}}$ and $\mathrm{ub}(f)_r$ for random instances of maximum cut. Finally, we consider some polynomial maximization problems on $[-1, 1]^n$ coming from small instances of MaxCut. An instance of MaxCut with vertex set $[n]$ and edge weights $w_{ij} \geq 0$ can be written as:

$$\mathrm{opt} := \max_{\mathbf{x} \in [-1,1]^n} f(\mathbf{x}), \text{ where } f(\mathbf{x}) := \frac{1}{4} \sum_{i,j \in [n]} w_{ij} (\mathbf{x}_i - \mathbf{x}_j)^2. \tag{5.10}$$

Note that while f is usually maximized over the discrete cube $\{-1, 1\}^n$, the formulation (5.10) is equivalent as f is convex.

Following [**Las11**], we create our instances by setting $w_{ij} = 0$ with probability p, and sampling w_{ij} uniformly from $[0, 1]$ otherwise. In Table 5.2, we list values of $\mathrm{ub}(f)_r^{\mathrm{pf}}$ and $\mathrm{ub}(f)_r$ for a few such random instances with $p = 1/2$ and $n = 8$. In each case, $\mathrm{ub}(f)_r^{\mathrm{pf}}$ provides a better bound than $\mathrm{ub}(f)_r$. In Table 5.3, we list the average over 50 randomly generated instances of the ratios:

$$\mathrm{Error} = \frac{\mathrm{opt} - \mathrm{ub}(f)_r}{\mathrm{opt}} \quad \text{and} \quad \mathrm{Error}_{\mathrm{pf}} = \frac{\mathrm{opt} - \mathrm{ub}(f)_r^{\mathrm{pf}}}{\mathrm{opt}}$$

for $r \leq 4$ and $p \in \{1/4,\ 1/2,\ 3/4\}$. Although it seems $\mathrm{ub}(f)_r^{\mathrm{pf}}$ is more sensitive to changes in the density of the instances, we find again that it provides a better bound in general than $\mathrm{ub}(f)_r$.

5.3.3. Comparison of $\mathrm{ub}(f, \mathbf{X}, \mu)_r$ for different sets X and measures μ. Finally, we compare the behaviour of the error $\mathrm{Error}(f; \mathbf{X}, \mu)_r$ for the test functions of Table 5.1 on different sets $\mathbf{X}$; namely the hypercube, the unit ball, and a regular octagon in $\mathbb{R}^2$. On the unit ball and the regular octagon, we consider the Lebesgue measure. On the hypercube, we consider both the Lebesgue measure and the Chebyshev measure. In each case, we compute the Lasserre bounds of order r in the range $1 \leq r \leq 20$, corresponding to sos-densities of degree up to 40. For the hypercube, simplex and unit

		$r=1$	$r=2$	$r=3$	$r=4$	$r=5$	$r=6$	opt
Ex1	$\mathrm{ub}(f)_r$	1.58	2.06	2.45	2.82	3.16	3.46	**6.09**
	$\mathrm{ub}(f)_r^{\mathrm{pf}}$	1.98	2.65	3.20	3.67	4.06	4.38	
Ex2	$\mathrm{ub}(f)_r$	2.20	2.77	3.27	3.73	4.16	4.54	**8.07**
	$\mathrm{ub}(f)_r^{\mathrm{pf}}$	2.61	3.41	4.11	4.72	5.25	5.71	
Ex3	$\mathrm{ub}(f)_r$	2.03	2.56	3.02	3.43	3.81	4.14	**7.24**
	$\mathrm{ub}(f)_r^{\mathrm{pf}}$	2.46	3.19	3.81	4.34	4.79	5.15	
Ex4	$\mathrm{ub}(f)_r$	1.59	2.05	2.44	2.81	3.13	3.42	**5.80**
	$\mathrm{ub}(f)_r^{\mathrm{pf}}$	1.98	2.62	3.15	3.60	3.97	4.28	

TABLE 5.2. Values of $\mathrm{ub}(f)_r$ and $\mathrm{ub}(f)_r^{\mathrm{pf}}$ for randomly generated instances of MAXCUT ($n = 8$, $p = 1/2$).

		$r=1$	$r=2$	$r=3$	$r=4$
$p=1/4$	Error	0.72	0.65	0.59	0.53
	$\mathrm{Error}_{\mathrm{pf}}$	0.66	0.57	0.48	0.40
$p=1/2$	Error	0.73	0.65	0.59	0.53
	$\mathrm{Error}_{\mathrm{pf}}$	0.68	0.56	0.47	0.39
$p=3/4$	Error	0.73	0.64	0.57	0.49
	$\mathrm{Error}_{\mathrm{pf}}$	0.65	0.53	0.43	0.35

TABLE 5.3. Average performance of the bounds $\mathrm{ub}(f)_r$ and $\mathrm{ub}(f)_r^{\mathrm{pf}}$ for random instances of MAXCUT ($n = 8$).

ball, closed form expressions for these moments are known (see, e.g., Table 1 in [**dKL19**]). For the octagon, they can be computed by triangulation.

5.3.4. The linear case. We consider first the case of a linear polynomial $f(\mathbf{x}) = f_{\mathrm{li}}(\mathbf{x}) = \mathbf{x}_1$ on $\mathbf{X} = [-1,1]^2$, equipped with the Lebesgue measure. Figure 5.5 shows the values of the parameters $\mathrm{Error}(f_{\mathrm{li}};\mathbf{X})_r$ and $\mathrm{Error}(f_{\mathrm{li}};\mathbf{X})_r \cdot r^2$. In accordance with Theorem 3.17 (and 3.16(ii)), it appears indeed that $\mathrm{Error}(f_{\mathrm{li}};\mathbf{X})_r = O(1/r^2)$, as suggested by the fact that the parameter $\mathrm{Error}(f_{\mathrm{li}};\mathbf{X})_r \cdot r^2$ approaches a constant value as r grows.

The unit ball. Next, we consider the unit ball B^2, again equipped with the Lebesgue measure. Figure 5.6 shows the values of the ratio:

$$\mathrm{Error}(f_*; B^2)_r / \mathrm{Error}(f_*; [-1,1]^2)_r \tag{5.11}$$

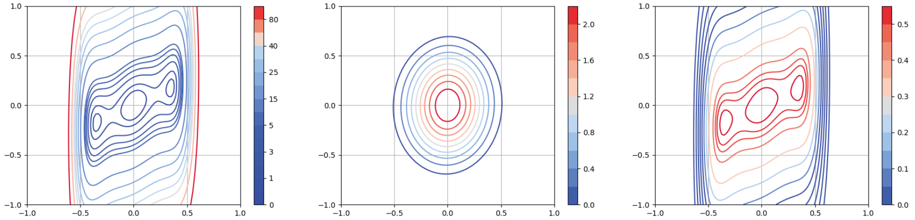

FIGURE 5.1. The Camel function (left) and its sum-of-squares densities corresponding to $\mathrm{ub}(f)_6$ (middle) and $\mathrm{ub}(f)_6^{\mathrm{pf}}$ (right) on the unit box.

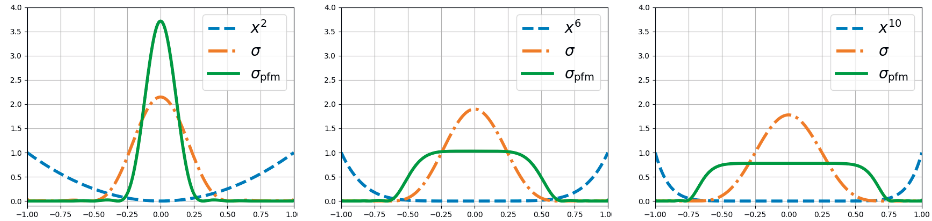

FIGURE 5.2. The functions $f(x) = x^{2k}$ and their sum-of-squares densities corresponding to $\mathrm{ub}(f)_6$ and $\mathrm{ub}(f)_6^{\mathrm{pf}}$ on the interval $[-1, 1]$ for $k = 1$ (left), $k = 3$ (middle) and $k = 5$ (right).

for $* \in \{\mathrm{li}, \mathrm{qu}, \mathrm{bo}, \mathrm{ma}, \mathrm{ca}, \mathrm{mo}\}$. Here, $f_{\mathrm{qu}} \in \mathbb{R}[\mathbf{x}]$ is the quadratic polynomial:

$$f_{\mathrm{qu}}(\mathbf{x}) = \mathbf{x}_1 + \mathbf{x}_2^2.$$

In each case, the ratio (5.11) appears to tend to a constant value, suggesting that the error $\mathrm{Error}(f_*; \mathbf{X})_r$ has similar asymptotic behaviour for $\mathbf{X} = [-1, 1]^2$ and $\mathbf{X} = B^2$. This matches the result of Theorem 3.19 both in the case of a minimizer on the boundary ($* \in \{\mathrm{li}, \mathrm{qu}\}$) and in the case of a minimizer in the interior ($* \in \{\mathrm{bo}, \mathrm{ma}, \mathrm{ca}, \mathrm{mo}\}$).

5.3.5. A regular octagon. Consider now a regular octagon:

$$\mathrm{Oct} = \mathrm{conv}\{(\pm 1, 0), (0, \pm 1), (\pm\frac{1}{2}\sqrt{2}, \pm\frac{1}{2}\sqrt{2})\} \subseteq [-1, 1]^2, \qquad (5.12)$$

equipped with the Lebesgue measure. This is an example of a convex body that is not ball-like (see Definition 3.22). Note that as a result, the strongest theoretical guarantee we have shown for the convergence rate of the Lasserre bounds on Oct is in $O(\log^2 r / r^2)$ (see Theorem 4.2). Figure 5.7 shows the

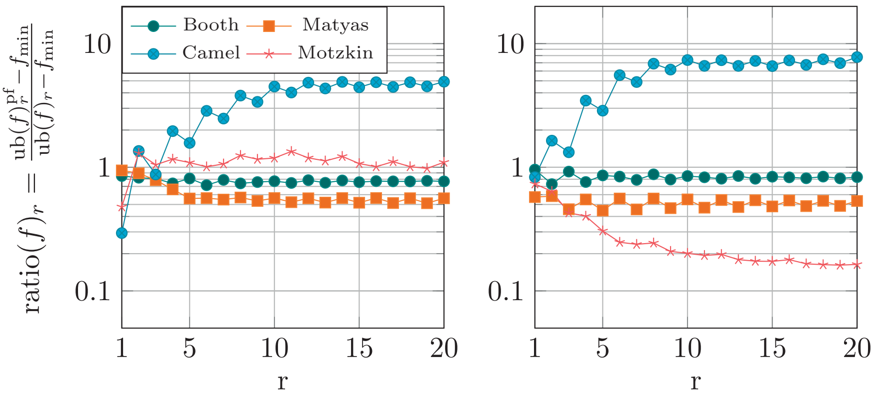

FIGURE 5.3. Comparison of the bounds $\text{ub}(f)_r$ and $\text{ub}(f)_r^{\text{pf}}$ for the four functions in Table 5.1, computed on the unit box (left) and unit ball (right).

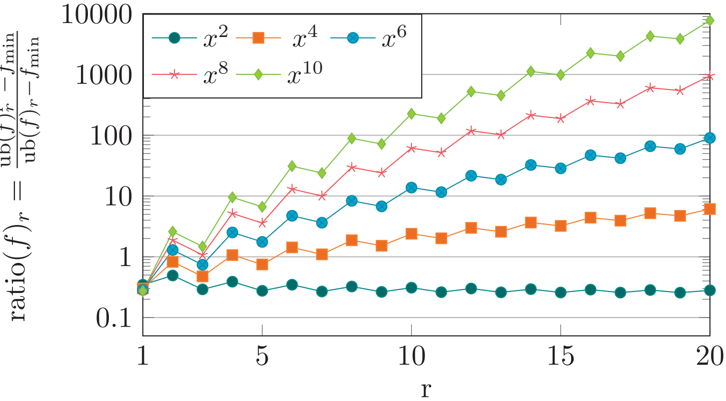

FIGURE 5.4. Comparison of the bounds $\text{ub}(f)_r$ and $\text{ub}(f)_r^{\text{pf}}$ for functions of the form $f(x) = x^{2k}$ on the interval $[-1, 1]$.

values of the ratio:

$$\text{Error}(f_*; \text{Oct})_r / \text{Error}(f_*; [-1, 1]^2)_r \tag{5.13}$$

for $* \in \{\text{li}, \text{qu}, \text{bo}, \text{ma}, \text{ca}, \text{mo}\}$. As for the unit ball, the ratio (5.13) seemingly tends to a constant value for each of the test polynomials. This indicates a similar asymptotic behaviour of the error $\text{Error}(f_*; \mathbf{X})_r$ for $\mathbf{X} = [-1, 1]^2$ and $\mathbf{X} = \text{Oct}$, and suggests that the convergence rate guaranteed by Theorem 4.2 might not be tight in this instance.

5.3.6. The Chebyshev measure. Finally, we consider the Chebyshev measure $d\mu(\mathbf{x}) = (1 - \mathbf{x}_1^2)^{-1/2}(1 - \mathbf{x}_2^2)^{-1/2} d\mathbf{x}$ on $[-1, 1]^2$, which we compare

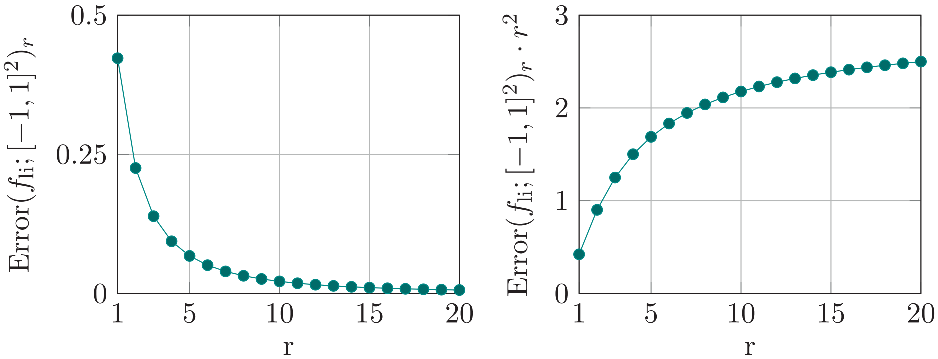

FIGURE 5.5. The error of the upper bounds $\mathrm{ub}(f)_r$ for $f(\mathbf{x}) = \mathbf{x}_1$ computed on $[-1,1]^2$ w.r.t. the Lebesgue measure.

to the Lebesgue measure. Figure 5.8 shows the values of the fraction:

$$\mathrm{Error}(f_*; [-1,1]^2, \mu)_r / \mathrm{Error}(f_*; [-1,1]^2)_r \tag{5.14}$$

for $* \in \{\mathrm{li}, \mathrm{qu}, \mathrm{bo}, \mathrm{ma}, \mathrm{ca}, \mathrm{mo}\}$. Again, we observe that the fraction (5.14) appears to tend to a constant value in each case, matching the result of Theorem 3.17.

Acknowledgements. We thank Utz-Uwe Haus and Angelika Wiegele for several discussions and their valuable suggestions. We also thank HPE Cray for making available their high-performance computing architecture.

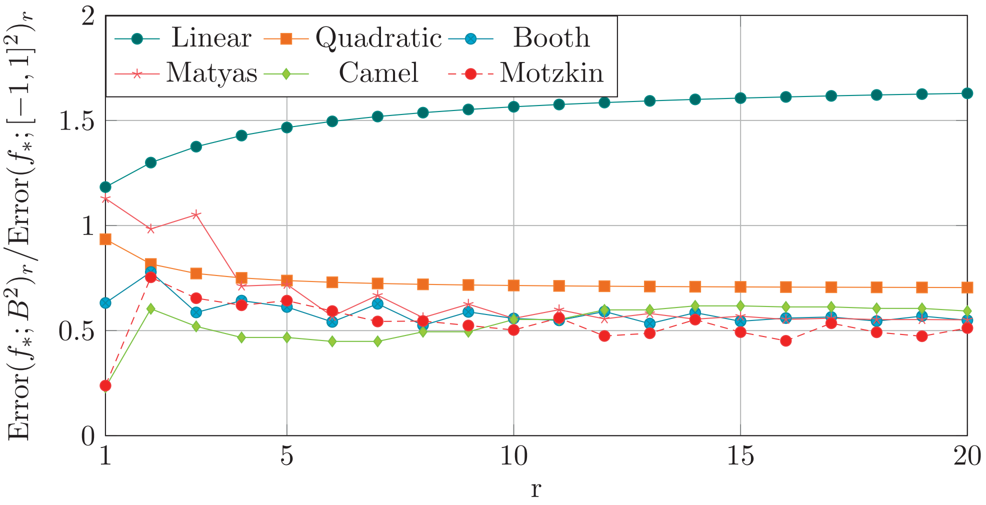

FIGURE 5.6. Comparison of the errors of the upper bounds $\mathrm{ub}(f)_r$ for the functions in Table 5.1 computed on $[-1,1]^2$ and the unit ball B^2 w.r.t. the Lebesgue measure.

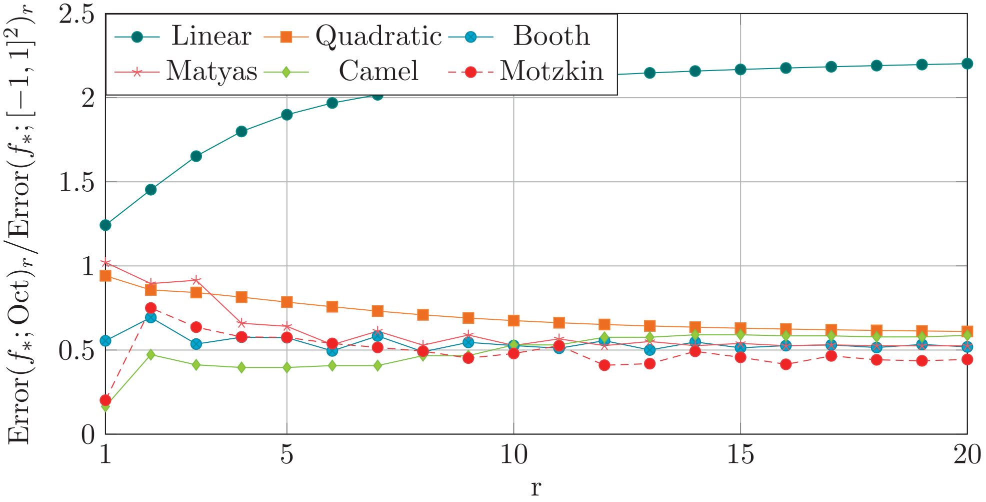

FIGURE 5.7. Comparison of the errors of the upper bounds $\mathrm{ub}(f)_r$ for the functions in Table 5.1 computed on $[-1,1]^2$ and the regular octagon Oct (see (5.12)) w.r.t. the Lebesgue measure.

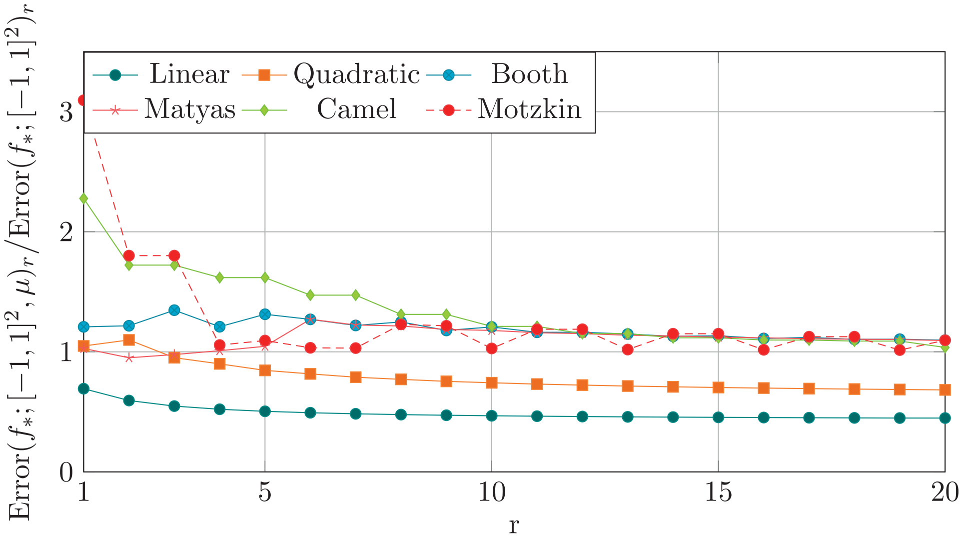

FIGURE 5.8. Comparison of the errors of the upper bounds $\mathrm{ub}(f)_r$ for the functions in Table 5.1 computed on $[-1,1]^2$ w.r.t. the Lebesgue and Chebyshev measures.

Part 2

Polynomial kernels and sum-of-squares hierarchies

CHAPTER 6

The polynomial kernel method

One of these days, I'm gonna get organezized.

From Taxi Driver

This chapter is based in part on my work [**Slo22**].

Let f be a polynomial of degree d. We consider the polynomial optimization problem:

$$f_{\min} := \min_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x}),$$

where $\mathbf{X} \subseteq \mathbb{R}^n$ is a compact semialgebraic set, defined by polynomials $g_1, \ldots, g_m \in \mathbb{R}[\mathbf{x}]$ as:

$$\mathbf{X} = \{\mathbf{x} \in \mathbb{R}^n : g_j(\mathbf{x}) \geq 0 \quad (1 \leq j \leq m)\}.$$

Recall from Chapter 2 that the cone $\mathcal{P}_+(\mathbf{X})$ of nonnegative polynomials on $\mathbf{X}$ may be approximated by the *quadratic module* $\mathcal{Q}(\mathbf{X})$ and the *preordering* $\mathcal{T}(\mathbf{X})$ of $\mathbf{X}$:

$$\mathcal{Q}(\mathbf{X}) := \Big\{ \sum_{j=0}^{m} \sigma_j g_j : \sigma_j \in \Sigma[\mathbf{x}] \Big\} \qquad \text{(where } g_0 := 1\text{)},$$
$$\mathcal{T}(\mathbf{X}) := \Big\{ \sum_{J \subseteq [m]} \sigma_J g_J : \sigma_J \in \Sigma[\mathbf{x}] \Big\} \quad \text{(where } g_J := \prod_{j \in J} g_j\text{)}.$$

By considering degree-restricted versions of the quadratic module and preordering, we obtain corresponding hierarchies of *lower* bounds on the global minimum $f_{\min}$ of f on $\mathbf{X}$:

$$\mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r := \sup\{\lambda \in \mathbb{R} : f - \lambda \in \mathcal{Q}(\mathbf{X})_{2r}\}, \tag{6.1}$$
$$\mathrm{lb}(f, \mathcal{T}(\mathbf{X}))_r := \sup\{\lambda \in \mathbb{R} : f - \lambda \in \mathcal{T}(\mathbf{X})_{2r}\}, \tag{6.2}$$

as well as upper bounds:

$$\mathrm{ub}(f, \mathbf{X}, \mu)_r := \inf_{q \in \Sigma[\mathbf{x}]_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mu(\mathbf{x}) : \int_{\mathbf{X}} q(\mathbf{x}) d\mu(\mathbf{x}) = 1 \right\}, \tag{6.3}$$

$$\mathrm{ub}(f, \mathcal{Q}(\mathbf{X}), \mu)_r := \inf_{q \in \mathcal{Q}(\mathbf{X})_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mu(\mathbf{x}) : \int_{\mathbf{X}} q(\mathbf{x}) d\mu(\mathbf{x}) = 1 \right\}, \tag{6.4}$$

$$\mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r := \inf_{q \in \mathcal{T}(\mathbf{X})_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mu(\mathbf{x}) : \int_{\mathbf{X}} q(\mathbf{x}) d\mu(\mathbf{x}) = 1 \right\}. \tag{6.5}$$

Outline. In this chapter, we present a unified approach to establish convergence rates for the lower bounds on certain structured sets $\mathbf{X}$. Namely, we show a connection between the behaviour of the *lower* bounds defined in (6.1), (6.2), the *upper* bounds defined in (6.3), (6.4), (6.5), and the *Christoffel-Darboux kernel* introduced in Chapter 1. As we explain in Section 6.1, this connection relies on an application of the *polynomial kernel method*, employing a perturbed version of the Christoffel-Darboux kernel. Following this approach, we are able to show strong convergence results in upcoming chapters for the lower bounds on the binary cube $\{0,1\}^n$ (Chapter 7), the unit ball and the standard simplex (Chapter 8), and the unit box $[-1,1]^n$ (Chapter 9).

The Christoffel-Darboux kernel has recently seen increased interest in the context of (polynomial) optimization [**dCGHL21, LP19, MPW$^+$21, PPL21**]. Of particular relevance is the recent work of Lasserre [**Las21**], where a link is established between this kernel and the hierarchy of lower bounds (6.1) (although this link is entirely different from the one we present below).

Essentially as a side result of our proof technique, we also obtain convergence rates for the corresponding hierarchies of *upper* bounds in these settings. Indeed, as we explain in Section 6.2, our approach actually makes this rather elementary. As we will see, however, the obtained rates mostly do not improve upon existing results. The exception is the binary cube (see Chapter 7).

Finally, our approach allows us to nicely present some known and new connections between the sum-of-squares hierarchies and bounds for polynomial optimization problems based on *cubature rules*. This is the topic of Section 6.3.

6.1. The polynomial kernel method

Here, we present the technique we use to prove the main results of Chapter 7, Chapter 8 and Chapter 9. It is inspired by the strategy used to prove convergence results for the lower bounds (6.1) on the hypersphere in [**FF21**].

Let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree $d \in \mathbb{N}$. Assume for simplicity that we are interested in analyzing the Putinar-type bound $\mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r$ for the minimization of f over a semialgebraic set $\mathbf{X}$. We then wish to show for each $r \in \mathbb{N}$ that:

$$f - f_{\min} + \varepsilon \in \mathcal{Q}(\mathbf{X})_{2r}$$

for some small $\varepsilon = \varepsilon(r) > 0$. Up to translation and scaling, we may assume that $f_{\min} = 0$ and that $\|f\|_{\mathbf{X}} := \max_{\mathbf{x} \in \mathbf{X}} |f(x)| = 1$. Recall from Chapter 1 that $\mathcal{P}(\mathbf{X})$ denotes the space of polynomials on $\mathbf{X}$. Suppose that we are able to construct an (invertible) linear operator $\mathbf{K} : \mathcal{P}(\mathbf{X})_d \to \mathcal{P}(\mathbf{X})_d$ which satisfies the following three properties:

$$\mathbf{K}(1) = 1, \tag{P1}$$

$$\mathbf{K}p \in \mathcal{Q}(\mathbf{X})_{2r} \quad \text{for all } p \in \mathcal{P}_+(\mathbf{X})_d \tag{P2}$$

$$\max_{\mathbf{x}\in\mathbf{X}} |\mathbf{K}^{-1}f(\mathbf{x}) - f(\mathbf{x})| \leq \varepsilon. \tag{P3}$$

We claim that we then have $f + \varepsilon \in \mathcal{Q}(\mathbf{X})_{2r}$. Indeed, since f is nonnegative on $\mathbf{X}$ by assumption, we know that $f(\mathbf{x}) + \varepsilon \geq \varepsilon$ for $\mathbf{x} \in \mathbf{X}$. By properties (P1) and (P3), it follows that $\mathbf{K}^{-1}(f + \varepsilon) \in \mathcal{P}_+(\mathbf{X})$. Using property (P2), we may thus conclude that:

$$f + \varepsilon = \mathbf{K}\big(\mathbf{K}^{-1}(f + \varepsilon)\big) \in \mathcal{Q}(\mathbf{X})_{2r},$$

meaning that $f_{\min} - \mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r \leq \varepsilon$.

We may thus establish convergence rates for the Putinar-type bounds (6.1) by showing the existence (for each $r \in \mathbb{N}$) of an operator $\mathbf{K}$ which satisfies (P1), (P2) and (P3) with $\varepsilon = \varepsilon(r)$ small. We summarize this observation in the following Lemma for future reference.

LEMMA 6.1. *Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact semialgebraic set and let f be a polynomial on $\mathbf{X}$ of degree d. Suppose that there exists a nonsingular linear operator $\mathbf{K} : \mathcal{P}(\mathbf{X})_d \to \mathcal{P}(\mathbf{X})_d$ which satisfies the properties* (P1), (P2) *and* (P3) *for certain $\varepsilon \geq 0$. Then $f_{\min} - \mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r \leq \varepsilon$.*

Note that by replacing the quadratic module $\mathcal{Q}(\mathbf{X})$ by the preordering $\mathcal{T}(\mathbf{X})$ in the above, we may use the exact same technique to establish convergence rates for the Schmüdgen-type bounds (6.2). For clarity, we will use the quadratic module throughout this chapter, but all results carry over immediately.

6.1.1. Constructing a linear operator. In light of Lemma 6.1 above, we wish to construct linear operators that satisfy (P1), (P2) and (P3). For this purpose, we apply the *polynomial kernel method.* Let us recall the setup for kernel operators from Chapter 1.

Let $\mathrm{K} : \mathbf{X} \times \mathbf{X} \to \mathbb{R}$ be a polynomial kernel on $\mathbf{X}$, meaning that $\mathrm{K}(\mathbf{x}, \mathbf{y})$ is a polynomial in the variables $\mathbf{x}, \mathbf{y}$. After choosing a finite, Borel measure μ supported on $\mathbf{X}$, we have an inner product $\langle \cdot, \cdot \rangle_\mu$ on the space $\mathcal{P}(\mathbf{X})$ of polynomials on $\mathbf{X}$ given by:

$$\langle p, q \rangle_\mu := \int_{\mathbf{X}} p(\mathbf{x})q(\mathbf{x})d\mu(\mathbf{x}) \quad (p, q \in \mathcal{P}(\mathbf{X})).$$

Using this inner product, we may associate a linear operator $\mathbf{K} : \mathcal{P}(\mathbf{X}) \to \mathcal{P}(\mathbf{X})$ to K by setting:

$$\mathbf{K}p(\mathbf{x}) := \langle \mathrm{K}(\mathbf{x}, \cdot), p \rangle_\mu = \int_{\mathbf{X}} \mathrm{K}(\mathbf{x}, \mathbf{y})p(\mathbf{y})d\mu(\mathbf{y}) \quad (p \in \mathbb{R}[\mathbf{x}]). \tag{6.6}$$

It turns out the operator $\mathbf{K}$ associated to a kernel K via (6.6) satisfies (P2) if the polynomial $\mathbf{x} \mapsto \mathrm{K}(\mathbf{x}, \mathbf{y})$ lies in $\mathcal{Q}(\mathbf{X})_{2r}$ for all *fixed* $\mathbf{y} \in \mathbf{X}$.

LEMMA 6.2. *Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact semialgebraic set, and let μ be a finite measure supported on $\mathbf{X}$. Let $Q \subseteq \mathbb{R}[\mathbf{x}]$ be a convex cone, and suppose that $\mathrm{K} : \mathbf{X} \times \mathbf{X} \to \mathbb{R}$ is a polynomial kernel for which $\mathrm{K}(\cdot, \mathbf{y}) \in Q$ for each $\mathbf{y} \in \mathbf{X}$ fixed. Then if $p \in \mathbb{R}[\mathbf{x}]$ is nonnegative on $\mathbf{X}$, we have $\mathbf{K}p \in Q$. That is, when selecting $Q = \mathcal{Q}(\mathbf{X})_{2r}$, the operator $\mathbf{K}$ associated to* K *satisfies* (P2).

PROOF. Let $\{(\mathbf{y}_i, w_i) : 1 \leq i \leq N\} \subseteq \mathbf{X} \times \mathbb{R}_{>0}$ be a cubature rule (see also Section 6.3 below) for the integration of polynomials of degree up to $\deg(p) + \deg(\mathrm{K})$ over $\mathbf{X}$ w.r.t. the measure μ, whose existence is guaranteed by Tchakaloff's Theorem [**Tch57**] (see also [**dKL19**]). Then by definition, we have:

$$\mathbf{K}p(\mathbf{x}) = \int_{\mathbf{X}} \mathrm{K}(\mathbf{x}, \mathbf{y}) p(\mathbf{y}) d\mu(\mathbf{y}) = \sum_{i=1}^{N} \mathrm{K}(\mathbf{x}, \mathbf{y}_i) w_i p(\mathbf{y}_i) \quad (\mathbf{x} \in \mathbf{X}).$$

As $w_i p(\mathbf{y}_i) \geq 0$ and $\mathrm{K}(\cdot, \mathbf{y}_i) \in Q$ for all $1 \leq i \leq N$, this shows that $\mathbf{K}p \in Q$. $\square$

Recall the *Christoffel-Darboux* kernel CD_{2r} of degree $2r$, which is defined in Chapter 1 in terms of a (graded) orthornormal basis $\{P_\alpha : \alpha \in \mathbb{N}^n\}$ for $\mathcal{P}(\mathbf{X})$ w.r.t. $\langle \cdot, \cdot \rangle_\mu$ as:

$$\mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}) := \sum_{|\alpha| \leq 2r} P_\alpha(\mathbf{x}) P_\alpha(\mathbf{y}).$$

The operator $\mathbf{CD}_{2r}$ associated to CD_{2r} via (6.6) *reproduces* the space of polynomials of degree up to $2r$, i.e., it satisfies:

$$\mathbf{CD}_{2r} p(\mathbf{x}) = p(\mathbf{x}) \quad (\mathbf{x} \in \mathbf{X}, \quad p \in \mathbb{R}[\mathbf{x}]_{2r}).$$

In other words, it is diagonal w.r.t. the basis $\{P_\alpha\}$, and its eigenvalues are all equal to 1. The idea now is to choose our kernel K by *perturbing* the Christoffel-Darboux kernel, that is, to consider a kernel of the form:

$$\mathrm{K}(\mathbf{x}, \mathbf{y}) = \mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda) := \sum_{|\alpha| \leq 2r} \lambda_\alpha P_\alpha(\mathbf{x}) P_\alpha(\mathbf{y}) \quad (\lambda_\alpha \in \mathbb{R}), \tag{6.7}$$

whose associated operator $\mathbf{K}$ has eigenvalues equal to the coefficients λ_α. It then remains to select these coefficients in such a way that the properties (P1), (P2) and (P3) are all satisfied.

6.1.2. Choosing the coefficients. Consider a kernel $\mathrm{K}(\mathbf{x}, \mathbf{y}) = \mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda)$, together with its associated operator $\mathbf{K}$. As we have:

$$\mathbf{K}(1) = \lambda_0,$$

it is clear that $\mathbf{K}$ satisfies (P1) if and only if $\lambda_0 = 1$. As we explain now, (P3) is satisfied when the coefficients λ_α, $|\alpha| \leq d$ are sufficiently close to 1. Recall that we consider in (P3) a polynomial f on $\mathbf{X}$ of degree d, whose sup-norm $\|f\|_{\mathbf{X}}$ over $\mathbf{X}$ is at most 1 by assumption, and that we wish to bound

$\|\mathbf{K}^{-1}f - f\|_{\mathbf{X}}$. As the polynomials $\{P_\alpha : |\alpha| \leq d\}$ form a basis of $\mathcal{P}(\mathbf{X})_d$, we may write:

$$f(\mathbf{x}) = \sum_{|\alpha| \leq d} f_\alpha(\mathbf{x}),$$

where $f_\alpha \in \mathrm{span}\{P_\alpha\}$. Assuming that $\lambda_0 = 1$ and $\lambda_\alpha \neq 0$ for $|\alpha| \leq d$, we then have $\mathbf{K}^{-1}f = \sum_{k=0}^{d}(1 - 1/\lambda_\alpha)f_\alpha$ and so:

$$\|\mathbf{K}^{-1}f - f\|_{\mathbf{X}} = \| \sum_{|\alpha|\leq d} (1 - 1/\lambda_\alpha)f_\alpha\|_{\mathbf{X}} \leq \max_{|\alpha|\leq d} \|f_\alpha\|_{\mathbf{X}} \cdot \sum_{|\alpha|\leq d} |1 - 1/\lambda_\alpha|. \quad (6.8)$$

In light of (6.8) and Lemma 6.2, we thus want to find coefficients $\lambda = (\lambda_\alpha)_{|\alpha|\leq 2r}$ such that:

(1) $\lambda_0 = 1$ and $\lambda_\alpha \approx 1$ for all $|\alpha| \leq d$;
(2) $x \mapsto \mathrm{K}(\mathbf{x}, \mathbf{y}) = \mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda) \in \mathcal{Q}(\mathbf{X})_{2r}$ for all $\mathbf{y} \in \mathbf{X}$.

How one might find λ simultaneously satisfying these both conditions depends on the structure of $\mathbf{X}$. Roughly speaking, one may distinguish three cases: a) the binary cube and unit sphere; b) the unit ball and the standard simplex; and c) the unit box $[-1, 1]^n$. These cases are discussed in detail in their respective chapters. Here, we sketch the main ideas.

The binary cube and unit sphere. On the binary cube and the unit sphere, the Christoffel-Darboux kernel may be written in a particularly simple form using addition formulas for *Krawtchouk* and *Gegenbauer* polynomials, respectively. Let us use the unit sphere to illustrate. As we saw already in Chapter 1, the Christoffel-Darboux kernel there is of the form:

$$\mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}) = \sum_{k=0}^{2r} \mathcal{G}_k^{(\frac{n-3}{2})}(\mathbf{x} \cdot \mathbf{y}). \quad (6.9)$$

Therefore, we should define our kernel K by:

$$\mathrm{K}(\mathbf{x}, \mathbf{y}) = \mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda) = \sum_{k=0}^{2r} \lambda_k \mathcal{G}_k^{(\frac{n-3}{2})}(\mathbf{x} \cdot \mathbf{y}).$$

Note that with respect to (6.7), we write here $\lambda_{|\alpha|} = \lambda_\alpha = \lambda_\beta$ whenever $|\alpha| = |\beta|$. It is not hard to see that $\mathrm{K}(\cdot, \mathbf{y}) \in \mathcal{Q}(\mathbf{X})_{2r}$ when the *univariate* polynomial:

$$q(x) := \sum_{k=0}^{2r} \lambda_k \mathcal{G}_k^{(\frac{n-3}{2})}(x)$$

is a sum of squares of degree $2r$. It thus remains to find a such a univariate $q \in \Sigma[x]_r$ for which $\lambda_0 = 1$ and $\lambda_k \approx 1$ for all $0 \leq k \leq d$. This is precisely what Fang and Fawzi do in [**FF21**], and what we will do (in the context of the binary cube) in Chapter 7. As we explain in more detail there, the problem of finding an optimal q in fact reduces to analyzing a univariate instance of the upper bounds (6.3). This reduction relies on the fact that the Gegenbauer polynomials form an orthogonal basis for $\mathbb{R}[x]$, and the coefficients λ_k are thus given (up to scaling) by:

$$\lambda_k \propto \int_{-1}^{1} \mathcal{G}_k^{(\frac{n-3}{2})}(x) q(x) (1-x^2)^{\frac{n-3}{2}} dx.$$

(Note in particular that the condition $\lambda_0 = 1$ means that q is a density!). As we have seen in Chapter 2 and Chapter 3, the upper bounds may be analyzed in the univariate case by considering the *roots* of the respective orthogonal polynomials. As these roots are well-understood for Gegenbauer and Krawtchouk polynomials, this yields the convergence rate $O(1/r^2)$ for the lower bounds $\mathrm{lb}(f, \mathcal{Q}(S^{n-1}))_r$ in [**FF21**], and the convergence rate of Theorem 7.1 for $\mathrm{lb}(f, \mathcal{Q}(\{0,1\}^n))_r$, respectively.

The unit ball and the standard simplex. For the unit ball B^n and the standard simplex Δ^n, expressions of the Christoffel-Darboux kernel as nice as (6.9) are not available. Fortunately, we can make use there of alternative expressions (8.17), (8.19) due to Xu [**Xu99, Xu98**]. These expressions are more complicated than (6.9), but they retain a key property; namely they allow $\mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda)$ to be written as the composition of a (relatively simple) multivariate polynomial with a univariate polynomial q of degree $2r$. The coefficients λ are then again related to the decomposition of q into the basis of Gegenbauer polynomials. Although it will be slightly more involved to show this, it turns out again that $\mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda)$ lies in $\mathcal{Q}(B^n)_{2r}$ (resp. $\mathcal{T}(\Delta^n)_{2r}$) if q is a sum of squares of degree $2r$. From there, the rest of the analysis is very similar to the case of the unit sphere. In Chapter 8, this yields convergence rates in $O(1/r^2)$ for the lower bounds $\mathrm{lb}(f, \mathcal{Q}(B^n))_r$ and $\mathrm{lb}(f, \mathcal{T}(\Delta^n))_r$.

The unit box $[-1,1]^n$. Finally, we treat the unit box $[-1,1]^n$ in Chapter 9. Our strategy there is a bit different. Namely, we make use of the so-called *Jackson kernel* $\mathrm{K}_{2r}^{\mathrm{jac}}$ on the interval $[-1,1]$. This kernel is equal to the perturbed Christoffel-Darboux kernel $\mathrm{CD}_{2r}(x, y; \lambda)$ on $[-1,1]$ with respect to the Chebyshev measure μ for a certain choice of coefficients λ. These coefficients are known to tend to 1 at a rate in $O(1/r^2)$. As mentioned in Chapter 2, this fact was already exploited in [**dKHL17**] to obtain convergence rates for the *upper* bounds $\mathrm{ub}(f, \mathcal{T}([-1,1]^n), \mu)_r$. Furthermore, it is known that $\mathrm{K}_{2r}^{\mathrm{jac}}(x, y) \geq 0$ for all $x, y \in [-1,1]$, which immediately implies that $\mathrm{K}_{2r}^{\mathrm{jac}}(\cdot, y)$ has a representation in $\mathcal{T}([-1,1])_{2r}$ (see Section 9.2). As we will see in Chapter 9, in order to move from the univariate to the multivariate case, one should then consider the kernel:

$$\mathrm{K}_{2r}^{\mathrm{jac}}(\mathbf{x}, \mathbf{y}) := \prod_{i=1}^{n} \mathrm{K}_{2r}^{\mathrm{jac}}(\mathbf{x}_i, \mathbf{y}_i).$$

6.1.3. The harmonic constant. We have so far conveniently ignored the constants $\|f_\alpha\|_{\mathbf{X}}$ that occur in (6.8). Using only the equivalence of norms on finite-dimensional vector spaces, one may say immediately that there exists a constant (depending on n and d) so that $\|f_\alpha\|_{\mathbf{X}} \leq c\|f\|_{\mathbf{X}}$ for all polynomials f of degree d. In the special cases we consider, however, one can say more. We will go into this in quite some detail in upcoming chapters. Roughly speaking, one may show for the unit sphere and the binary cube that these constants may be bounded *independently* of the dimension n, see [**FF21**] and Section 7.4, respectively. For the unit ball, the standard simplex and the box $[-1, 1]^n$, the constants may be shown to depend polynomially on n (for fixed d) and polynomially on d (for fixed n), see Section 8.3.1 and Section 9.3, respectively. As we discuss later, it is an open question whether a bound independent of the dimension may be shown in these cases as well.

6.2. Analysis of the hierarchies of upper bounds

In this section, we explain how to use the polynomial kernel method developed in this chapter to obtain convergence results for the *upper bounds* (6.4), (6.5). As we have seen in Chapter 4, one method of analyzing the behaviour of the upper bounds (employed, e.g., in [**dKHL17, dKL10, dKL18, dKLS17, SL20**]) is to exhibit an explicit probability density $\sigma \in \mathbb{R}[\mathbf{x}]$ on $(\mathbf{X}, \mu)$ which lies in the appropriate cone, and for which the difference:

$$\int_{\mathbf{X}} f(\mathbf{x})\sigma(\mathbf{x})d\mu(\mathbf{x}) - f_{\min}$$

can be bounded from above. We exhibit such a σ here based on the perturbed Christoffel-Darboux kernels (6.7) constructed in Section 6.1.

Write $\mathrm{K}(\mathbf{x}, \mathbf{y}) = \mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda)$ for such a kernel. Assume that the associated operator $\mathbf{K}$ satisfies (P1) and that:

$$\max_{\mathbf{x}\in\mathbf{X}} |\mathbf{K}f(\mathbf{x}) - f(\mathbf{x})| \leq \varepsilon. \tag{6.10}$$

for some $\varepsilon > 0$. The difference with (P3) is that we consider there the inverse kernel operator $\mathbf{K}^{-1}$. As we explained though, we establish (P3) by showing that the eigenvalues λ of $\mathbf{K}$ are sufficiently close to 1, in which case (6.10) also holds. For details, see Section 7.2.

Assume finally that $\mathrm{K}(\cdot, \mathbf{y}) \in \mathcal{Q}(\mathbf{X})_{2r}$ for all $\mathbf{y} \in \mathbf{X}$ (which we recall is the condition we use to show that $\mathbf{K}$ satisfies (P2) in Lemma 6.2 above). Let $\mathbf{x}^* \in \mathbf{X}$ be a global minimizer of f over $\mathbf{X}$, and consider the polynomial σ given by:

$$\sigma(\mathbf{x}) = \mathrm{K}(\mathbf{x}, \mathbf{x}^*) \quad (\mathbf{x} \in \mathbf{X}).$$

By assumption, $\sigma \in \mathcal{Q}(\mathbf{X})_{2r}$. As $\mathbf{K}$ satisfies (P1), we have:

$$\int_{\mathbf{X}} \sigma(\mathbf{x})d\mu(\mathbf{x}) = \mathbf{K}(1) = 1,$$

meaning that σ is a probability density on $\mathbf{X}$ w.r.t. μ. The polynomial σ is thus a feasible solution to (6.4). Furthermore, by (6.10) we have:

$$\int_{\mathbf{X}} f(\mathbf{x})\sigma(\mathbf{x})d\mu(\mathbf{x}) - f_{\min} = \mathbf{K}f(\mathbf{x}^*) - f(\mathbf{x}^*) \leq \varepsilon.$$

We may thus conclude that $\mathrm{ub}(f, \mathcal{Q}(\mathbf{X}), \mu)_r \leq \varepsilon$. By replacing $\mathcal{Q}(\mathbf{X})$ by $\mathcal{T}(\mathbf{X})$ in the above, the same argument works for $\mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r$.

The upshot is that whenever we apply the polynomial kernel method in future chapters to obtain convergence rates for the hierarchies of lower bounds, we more or less automatically obtain rates for the corresponding *upper* bounds as well. As we have mentioned, these rates are usually not better than the ones yielded by a direct analysis.

6.3. Sum-of-squares hierarchies and cubature rules

To close this chapter, we present some connections between sum-of-squares hierarchies and cubature rules. Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact set equipped with a measure μ. Let $r \in \mathbb{N}$. We say that:

$$\mathcal{W} := \{(\mathbf{x}_j, w_j) : 1 \leq j \leq N\} \subseteq \mathbf{X} \times \mathbb{R} \tag{6.11}$$

is a *cubature rule* for $(\mathbf{X}, \mu)$ of power r and size N if:

$$\int_{\mathbf{X}} p(\mathbf{x})d\mu(\mathbf{x}) = \sum_{i=1}^{N} w_j p(\mathbf{x}_j)$$

for all polynomials $p \in \mathcal{P}(\mathbf{X})$ of degree r. We say that $\mathcal{W}$ is *positive* when $w_j > 0$ for all $1 \leq j \leq N$.

6.3.1. An upper bound. Let us move back to the context of minimizing a polynomial f of degree d over $\mathbf{X}$. Given a cubature rule $\mathcal{W}$ for $(\mathbf{X}, \mu)$ as in (6.11), one may naturally define an *upper* bound on the global minimum $f_{\min}$ of f by setting:

$$\mathrm{ub}(f, \mathbf{X}, \mathcal{W})_{\mathrm{cub}} := \min_{1 \leq j \leq N} f(\mathbf{x}_j) \geq f_{\min}.$$

Indeed, we have simply sampled the feasible region $\mathbf{X}$. The point is that when $\mathcal{W}$ is positive and has power at least $2r+d$, the resulting bound $\mathrm{ub}(f, \mathbf{X}, \mathcal{W})_{\mathrm{cub}}$ is at least as good as the measure-based upper bound $\mathrm{ub}(f, \mathbf{X}, \mu)_r$. This fact was first pointed out in [**MPSV20**]. It was used by de Klerk and Laurent [**dKL19**] to show tightness of their analysis in [**dKL20b**] of the measure-based bounds on $[-1, 1]^n$. Some further implications can be found in [**dKL19**] as well. It can be proven as follows. Let $\sigma^* \in \Sigma[\mathbf{x}]_r$ be an optimum solution to the program (6.3) defining $\mathrm{ub}(f, \mathbf{X}, \mu)_r$, meaning σ^* is a density and

$$\int_{\mathbf{X}} f(\mathbf{x})\sigma^*(\mathbf{x})d\mu(\mathbf{x}) = \mathrm{ub}(f, \mathbf{X}, \mu)_r.$$

Using the definition of a cubature rule, the fact that $\deg(\sigma^*) = 2r$ and the fact that $w_j > 0$, we see that:

$$\begin{aligned}\mathrm{ub}(f, \mathbf{X}, \mu)_r = \int_{\mathbf{X}} f(\mathbf{x})\sigma^*(\mathbf{x})d\mu(\mathbf{x}) &= \sum_{i=1}^{N} w_j f(\mathbf{x}_j)\sigma^*(\mathbf{x}_j) \\ &\geq \min_{1\leq j\leq N} f(\mathbf{x}_j) \cdot \sum_{i=1}^{N} w_j \sigma^*(\mathbf{x}_j) = \mathrm{ub}(f, \mathbf{X}, \mathcal{W})_{\mathrm{cub}}.\end{aligned}$$

A similar argument works for the bounds $\mathrm{ub}(f, \mathcal{Q}(\mathbf{X}), \mu)_r$, $\mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r$ and the push-forward bound $\mathrm{ub}(f, \mathbf{X}, \mu)_r$.

Optimizing over gridsets. Having seen that optimizing over the points of a cubature rule on $\mathbf{X}$ yields a good upper bound on $f_{\min}$, it is natural to ask what happens when we optimize instead over a simpler subset of $\mathbf{X}$, and indeed, this question has been considered in the literature. For example, optimizing over the mesh $G_r = \{i/r : 0 \leq i \leq r\}^n \subseteq [0,1]^n$, $r \in \mathbb{N}$ yields an upper bound which is within $O(1/r^2)$ of the true minimum of a polynomial f on $[0,1]^n$, see [**dKLLS17**]. An analogous statement holds on the simplex, see [**dKLVS17**]. These upper bounds thus have the same asymptotic error in r as Lasserre's measure-based upper bounds. Importantly, though, we should note that for fixed $r \in \mathbb{N}$, the set G_r (used for the hypercube) is of exponential size $(r+1)^n$, whereas the measure-based bounds may be computed in polynomial time.

6.3.2. A lower bound. One may also use cubature rules to define *lower* bounds on $f_{\min}$. As we will see now, though, this is a bit more involved. We begin with a slight reinterpretation of the polynomial kernel method above. Let K be a polynomial kernel on $\mathbf{X}$ whose associated operator $\mathbf{K}$ satisfies properties (P1) and (P2). Consider the parameter:

$$\mathrm{lb}(f, \mathbf{X}, \mathrm{K})_{\mathrm{harm}} := \min_{\mathbf{x}\in\mathbf{X}} \mathbf{K}^{-1} f(\mathbf{x}). \tag{6.12}$$

As the notation suggests, this parameter is a lower bound on $f_{\min}$. In fact, it is a lower bound on $\mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r$. Indeed, the function

$$\mathbf{K}^{-1} f - \mathrm{lb}(f, \mathbf{X}, \mathrm{K})_{\mathrm{harm}}$$

is nonnegative on $\mathbf{X}$ by definition. By (P1) and (P2), we thus have:

$$f - \mathrm{lb}(f, \mathbf{X}, \mathrm{K})_{\mathrm{harm}} = \mathbf{K}\big(\mathbf{K}^{-1} f - \mathrm{lb}(f, \mathbf{X}, \mathrm{K})_{\mathrm{harm}}\big) \in \mathcal{Q}(\mathbf{X})_{2r},$$

showing that (cf. (6.1)):

$$\mathrm{lb}(f, \mathbf{X}, \mathrm{K})_{\mathrm{harm}} \leq \mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r \leq f_{\min}. \tag{6.13}$$

This gives us the following (informal) reinterpretation of Lemma 6.1 (which we find rather clarifying).

Remark 6.3. *Let $\varepsilon > 0$, and suppose that we are able to find a kernel* K *on* $\mathbf{X}$ *which satisfies* (P1), (P2) *and for which:*

$$f_{\min} - \mathrm{lb}(f, \mathbf{X}, \mathrm{K})_{\mathrm{harm}} \leq \varepsilon. \tag{6.14}$$

Then $f_{\min} - \mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r \leq \varepsilon$.

In light of this remark, all that remains to analyze $\mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r$ is thus to construct kernels K for which $f_{\min} - \mathrm{lb}(f, \mathbf{X}, \mathrm{K})_{\mathrm{harm}}$ is small, which is precisely what we will do in upcoming chapters.

From a practical point of view, the problem with the bound $\mathrm{lb}(f, \mathbf{X}, \mathrm{K})_{\mathrm{harm}}$ is that it is not clear how one may compute it. Indeed, even if one is able to evaluate the function $\mathbf{K}^{-1} f(\mathbf{x})$ for all $\mathbf{x} \in \mathbf{X}$, one still has to determine the minimum (6.12). We can address this problem using a cubature rule; which is an idea due essentially to Christancho and Velasco [**CV22**] (who consider specifically the case $\mathbf{X} = S^{n-1}$). Let $\mathcal{W} = \{(\mathbf{x}_j, w_j) : 1 \leq j \leq N\}$ again be a positive cubature rule for $(\mathbf{X}, \mu)$ of power $2r + d$. Consider the parameter:

$$\mathrm{lb}(f, \mathbf{X}, \mathrm{K}, \mathcal{W})_{\mathrm{cub}} := \min_{1 \leq j \leq N} \mathbf{K}^{-1} f(\mathbf{x}_j),$$

which may be computed by evaluating $\mathbf{K}^{-1} f(\mathbf{x})$ at only N points $\mathbf{x} \in \mathbf{X}$. Again, as our notation suggests, this parameter is a lower bound on $f_{\min}$. More precisely, we have:

$$\mathrm{lb}(f, \mathbf{X}, \mathrm{K})_{\mathrm{harm}} \leq \mathrm{lb}(f, \mathbf{X}, \mathrm{K}, \mathcal{W})_{\mathrm{cub}} \leq \mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r \leq f_{\min}. \tag{6.15}$$

The first inequality is by definition. The second inequality follows in a way similar to (6.13) (the attentive reader may also recognize the argument from the proof of Lemma 6.2): Note that $\mathbf{K}^{-1} f(\mathbf{x}_j) - \mathrm{lb}(f, \mathbf{X}, \mathrm{K}, \mathcal{W})_{\mathrm{cub}} \geq 0$ for all $1 \leq j \leq N$ by definition. Assuming that $\mathrm{K}(\cdot, \mathbf{y}) \in \mathcal{Q}(\mathbf{X})_{2r}$ for all $\mathbf{y} \in \mathbf{X}$, it follows that:

$$\begin{aligned} f(\mathbf{x}) - \mathrm{lb}(f, \mathbf{X}, \mathrm{K}, \mathcal{W})_{\mathrm{cub}} &= \mathbf{K}\mathbf{K}^{-1}\big(f(\mathbf{x}) - \mathrm{lb}(f, \mathbf{X}, \mathrm{K}, \mathcal{W})_{\mathrm{cub}}\big) \\ &= \int_{\mathbf{X}} \mathrm{K}(\mathbf{x}, \mathbf{y})\big(\mathbf{K}^{-1} f(\mathbf{y}) - \mathrm{lb}(f, \mathbf{X}, \mathrm{K}, \mathcal{W})_{\mathrm{cub}}\big) d\mu(\mathbf{y}) \\ &= \sum_{j=1}^{N} \mathrm{K}(\mathbf{x}, \mathbf{x}_j) w_j \big(\mathbf{K}^{-1} f(\mathbf{x}_j) - \mathrm{lb}(f, \mathbf{X}, \mathrm{K}, \mathcal{W})_{\mathrm{cub}}\big) \end{aligned}$$

which shows that $f(\mathbf{x}) - \mathrm{lb}(f, \mathbf{X}, \mathrm{K}, \mathcal{W})_{\mathrm{cub}}$ lies in $\mathcal{Q}(\mathbf{X})_{2r}$, and thus that:

$$\mathrm{lb}(f, \mathbf{X}, \mathrm{K}, \mathcal{W})_{\mathrm{cub}} \leq \mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r.$$

In upcoming chapters and in the work [**FF21**] on S^{n-1}, convergence rates for the sum-of-squares bounds $\mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r$ are always shown by constructing kernels K_r in $\mathcal{Q}(\mathbf{X})_{2r}$ of increasing degree r for which the (weaker) bounds $\mathrm{lb}(f, \mathbf{X}, \mathrm{K}_r)_{\mathrm{harm}}$ are close to $f_{\min}$. By (6.15), this immediately implies convergence rates for the bounds $\mathrm{lb}(f, \mathbf{X}, \mathrm{K}_r, \mathcal{W})_{\mathrm{cub}}$ as well.

The hypersphere. Consider the special case of minimization of a polynomial f over the unit sphere S^{n-1}. In [**FF21**], the authors construct kernels K_r in $\mathcal{Q}(S^{n-1})_{2r}$ that satisfy (6.14) with $\varepsilon = O(1/r^2)$, thereby showing that the bounds $\mathrm{lb}(f, \mathcal{Q}(S^{n-1}))_r$ converge to $f_{\min}$ at a rate in $O(1/r^2)$. By (6.15), this directly implies the same convergence rate for the cubature based bounds $\mathrm{lb}(f, S^{n-1}, \mathrm{K}_r, \mathcal{W})_{\mathrm{cub}}$. In light of this fact, it would be interesting to actually compute these bounds. There are two potential obstacles in doing so. First, one needs to be able to evaluate $\mathbf{K}_r^{-1} f(\mathbf{x})$ for $\mathbf{x} \in \mathbf{X}$. Second, one needs a (preferably small) positive cubature rule $\mathcal{W}$ of power $2r + \deg(f)$ for S^{n-1} (w.r.t. the uniform measure). As the kernels K_r are of the form (6.7), the first obstacle can be overcome using the formula (6.9). In [**CV22**], the authors overcome the second obstacle in part by using numerically stable cubature rules of size $N \approx r^n$. Fixing n, this results in an alternative hierarchy of tractable lower bounds on $f_{\min}$ which may be computed without the need to solve any semidefinite programs.

6.4. Discussion

We have presented a unified approach to prove convergence rates for the lower bounds based on kernel operators, and the Christoffel-Darboux kernel in particular. In upcoming chapters, we specialize our method to particular sets $\mathbf{X}$, which include the binary cube, the unit ball, the standard simplex and the box $[-1, 1]^n$.

Connecting different hierarchies. Our method also reveals some connections between upper and lower bounds using sums of squares and bounds using cubature rules, which we have presented in Section 6.2 and Section 6.3. They may be summarized as follows. Let $\mathbf{X}$ be a compact semialgebraic set equipped with a measure μ. Suppose we have a kernel K on $\mathbf{X}$ in $\mathcal{Q}(\mathbf{X})_{2r}$ and assume that its associated operator $\mathbf{K}$ satisfies $\mathbf{K}(1) = 1$. Furthermore, let $\mathcal{W} = \{(\mathbf{x}_j, w_j) : 1 \leq j \leq N\}$ be a positive cubature rule for $(\mathbf{X}, \mu)$ of power $2r + \deg(f)$. Then we have:

$$
\begin{array}{ccccccccccccc}
\mathrm{lb}(f,\mathbf{X},\mathrm{K})_{\mathrm{harm}} & \leq & \mathrm{lb}(f,\mathbf{X},\mathrm{K},\mathcal{W})_{\mathrm{cub}} & \leq & \mathrm{lb}(f,\mathcal{Q}(\mathbf{X}))_r & \leq & f_{\min} & \leq & \mathrm{ub}(f,\mathbf{X},\mathcal{W})_{\mathrm{cub}} & \leq & \mathrm{ub}(f,\mathbf{X})_r & \leq & \mathbf{K}f(\mathbf{x}^*) \\
\shortparallel & & \shortparallel & & & & \shortparallel & & \shortparallel & & & & \rotatebox{-90}{$\leq$} \\
\min\limits_{\mathbf{x}\in\mathbf{X}} \mathbf{K}^{-1} f(\mathbf{x}) & & \min\limits_{1\leq j\leq N} \mathbf{K}^{-1} f(\mathbf{x}_j) & & & & f(\mathbf{x}^*) & & \min\limits_{1\leq j\leq N} f(\mathbf{x}_j) & & & & \min\limits_{\mathbf{x}\in\mathbf{X}} \mathbf{K}f(\mathbf{x})
\end{array}
$$

Possible extensions. In Part 1 of this thesis, we saw that one can show (near) tight convergence rates for the measure-based upper bounds in a very general setting. Obtaining such general results for the lower bounds using our method of proof seems difficult; so far it can only be applied to highly structured sets $\mathbf{X}$, where the Christoffel-Darboux kernel admits a nice expression. It is an important open question whether one may adapt our method to settings with less structure.

In regards to the lower bounds based on cubature rules discussed in Section 6.3, it should in principle be possible to perform explicit computations on some of the sets we consider in future chapters. In order to do so, one would have to adapt the ideas used in [**CV22**] for the hypersphere S^n. Most importantly, one would need to find good cubature rules for these sets.

Acknowledgments. We thank Monique Laurent and Mauricio Velasco for fruitful discussions on the presentation of the material and the connection to cubature rules.

CHAPTER 7

Application: The binary cube

I have completely forgotten the symbolic calculus.

Emily Noether

This chapter is based on my joint work [**SL21b**] *with Monique Laurent.*

We consider the problem of minimizing a polynomial $f \in \mathbb{R}[\mathbf{x}]$ of degree $d \leq n$ over the n-dimensional binary hypercube $\mathbb{B}^n = \{0,1\}^n$, i.e., of computing:

$$f_{\min} := \min_{\mathbf{x} \in \mathbb{B}^n} f(\mathbf{x}). \tag{7.1}$$

This optimization problem is NP-hard in general. Indeed, as is well-known, one can model an instance of MAX-CUT on the complete graph K_n with edge weights $w = (w_{ij})$ as a problem of the form (7.1) by setting:

$$f(\mathbf{x}) = -\sum_{1 \leq i < j \leq n} w_{ij}(\mathbf{x}_i - \mathbf{x}_j)^2.$$

As another example one can compute the stability number $\alpha(G)$ of a graph $G = (V, E)$ via the program

$$\alpha(G) = \max_{\mathbf{x} \in \mathbb{B}^{|V|}} \sum_{i \in V} \mathbf{x}_i - \sum_{\{i,j\} \in E} \mathbf{x}_i \mathbf{x}_j.$$

One may replace the binary cube $\mathbb{B}^n = \{0,1\}^n$ by the discrete cube $\{\pm 1\}^n$, in which case maximizing a quadratic polynomial $\mathbf{x}^\top A \mathbf{x}$ has many other applications, e.g., to MAX-CUT [**GW95**], to the cut norm [**AN04**], or to correlation clustering [**BBC04**]. Approximation algorithms are known depending on the structure of the matrix A (see [**AN04, CW04, GW95**]), but the problem is known to be NP-hard to approximate within any factor less than 13/11 [**ABH⁺05**].

Sum-of-squares hierarchies on the binary cube. The binary cube is a semialgebraic set, with description:

$$\mathbb{B}^n = \{\mathbf{x}_i^2 - \mathbf{x}_i = 0 : 1 \leq i \leq n\}.$$

The space $\mathcal{P}(\mathbb{B}^n)$ of polynomials on $\{0,1\}^n$ is given by $\mathbb{R}[\mathbf{x}]/\mathcal{I}$, where $\mathcal{I}$ is the ideal generated by the polynomials $\mathbf{x}_1 - \mathbf{x}_1^2, \ldots, \mathbf{x}_n - \mathbf{x}_n^2$. Alternatively, we have:

$$\mathcal{P}(\mathbb{B}^n) = \text{span}\{\mathbf{x}^\alpha : \alpha \in \{0,1\}^n\}.$$

Lower bounds. The quadratic module $\mathcal{Q}(\mathbb{B}^n)$ and preordering $\mathcal{T}(\mathbb{B}^n)$ in this case are rather special; namely, we have:

$$\mathcal{Q}(\mathbb{B}^n) = \mathcal{T}(\mathbb{B}^n) = \{p \in \mathbb{R}[\mathbf{x}] : p \text{ is a sum of squares on } \mathbb{B}^n\}.$$

Here, 'p is a sum of squares on $\mathbb{B}^n$' means that there exists a sum of squares $q \in \Sigma[\mathbf{x}]$ such that $p(\mathbf{x}) = q(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{B}^n$; or alternatively that $p - q$ lies in the ideal $\mathcal{I}$. The lower bounds $\text{lb}(f, \mathcal{Q}(\mathbb{B}^n))_r$ and $\text{lb}(f, \mathcal{T}(\mathbb{B}^n))_r$ on $f_{\min}$ are therefore equal, and given by:

$$\sup_{\lambda \in \mathbb{R}} \{f - \lambda \text{ is a sum-of-squares of degree at most } 2r \text{ on } \mathbb{B}^n\}.$$

For simplicity, we will write $\text{lb}(f)_r$ for these bounds throughout this chapter.

Another peculiarity of this setting is that the bounds $\text{lb}(f)_r$ have finite convergence: $\text{lb}(f)_r = f_{\min}$ for $r \geq n$ [**Las01, Lau03a**]. In fact, it has been shown in [**STKI17**] that the bound $\text{lb}(f)_r$ is exact already for $2r \geq n + d - 1$. That is:

$$\text{lb}(f)_r = f_{\min} \text{ for } r \geq \frac{n+d-1}{2}. \tag{7.2}$$

In addition, it is shown in [**STKI17**] that the bound $\text{lb}(f)_r$ is exact for $2r \geq n + d - 2$ when the polynomial f has only monomials of even degree. This extends an earlier result of [**FSP16**] shown for quadratic forms ($d = 2$), which applies in particular to the case of MAX-CUT. Furthermore, this result is tight for MAX-CUT, since one needs to go up to order $2r \geq n$ in order to reach finite convergence (in the cardinality case when all edge weights are 1) [**Lau03b**]. Similarly, the result (7.2) is tight when d is even and n is odd [**KLM16**].

Upper bounds. Let μ be the uniform probability measure on $\mathbb{B}^n$. In addition to the *lower* bound $\text{lb}(f)_r$, we also consider the measure-based *upper* bound $\text{ub}(f, \mathbb{B}^n, \mu)_r$ on $f_{\min}$, which we recall is defined as follows:

$$\text{ub}(f, \mathbf{X}, \mu)_r := \inf_{s \in \Sigma[\mathbf{x}]_{2r}} \left\{ \int_{\mathbb{B}^n} f(\mathbf{x}) \cdot s(\mathbf{x}) d\mu(\mathbf{x}) : \int_{\mathbb{B}^n} s(\mathbf{x}) d\mu(\mathbf{x}) = 1 \right\}. \tag{7.3}$$

Throughout this chapter, we shall simply write $\text{ub}(f)_r$ for this bound. The bound $\text{ub}(f)_r$ also converges to $f_{\min}$ in finitely many steps [**Las11**]; in fact it is not difficult to see that it is exact at order $r = n$ and that this is tight (see Section 7.3).

Outline. The main contribution of this chapter is an analysis of the quality of the bounds $\mathrm{lb}(f)_r$ for parameters $r, n \in \mathbb{N}$ with $2r < n+d-1$, i.e., for which the bounds are not exact. The following is our main result, which expresses the error of the bound $\mathrm{lb}(f)_r$ in terms of the roots of Krawtchouk polynomials, which we recall from Chapter 1 are classical univariate orthogonal polynomials with respect to a discrete measure on the set $\{0, 1, \ldots, n\}$ (see also Section 7.1 below).

THEOREM 7.1. *Fix $d \leq n$ and let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree d. For $r, n \in \mathbb{N}$, let ξ_r^n be the least root of the degree r Krawtchouk polynomial* (7.11) *with parameter n. Then, if $(r+1)/n \leq 1/2$ and $d(d+1) \cdot \xi_{r+1}^n/n \leq 1/2$, we have:*

$$\frac{f_{\min} - \mathrm{lb}(f)_r}{\|f\|_\infty} \leq 2C_d \cdot \xi_{r+1}^n/n. \tag{7.4}$$

Here $C_d > 0$ is an absolute constant depending only on d and we set $\|f\|_\infty := \max_{\mathbf{x}\in\mathbb{B}^n} |f(\mathbf{x})|$.

The extremal roots of Krawtchouk polynomials are well-studied in the literature. The following result of Levenshtein [**Lev98**] shows their asymptotic behaviour.

THEOREM 7.2 ([**Lev98**], Section 5). *For $t \in [0, 1/2]$, define the function*

$$\varphi(t) = 1/2 - \sqrt{t(1-t)}. \tag{7.5}$$

Then the least root ξ_r^n of the degree r Krawtchouk polynomial with parameter n satisfies

$$\xi_r^n/n \leq \varphi(r/n) + c \cdot (r/n)^{-1/6} \cdot n^{-2/3} \tag{7.6}$$

for some universal constant $c > 0$.

Applying (7.6) to (7.4), we find that the relative error of the bound $\mathrm{lb}(f)_r$ in the regime $r \approx t \cdot n$ behaves as the function $\varphi(t) = 1/2 - \sqrt{t(1-t)}$, up to a term in $O(1/n^{2/3})$, which vanishes as n tends to ∞. As an illustration, Figure 7.1 in Section 7.5 shows the function $\varphi(t)$.

As we explain below, we will use the polynomial kernel method of Chapter 6 to prove Theorem 7.1. As we saw in Section 6.2, we can show the following analog of Theorem 7.1 for the upper bounds $\mathrm{ub}(f)_r$ as a side result. To our knowledge this is the first analysis of the upper bounds on $\mathbb{B}^n$.

THEOREM 7.3. *Fix $d \leq n$ and let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree d. Then, for any $r, n \in \mathbb{N}$ with $(r+1)/n \leq 1/2$, we have:*

$$\frac{\mathrm{ub}(f)_r - f_{\min}}{\|f\|_\infty} \leq C_d \cdot \xi_{r+1}^n/n,$$

where $C_d > 0$ is the constant mentioned in Theorem 7.1.

Note that the above analysis of $\mathrm{ub}(f)_r$ does not require any condition on the size of ξ^n_{r+1} as was necessary for the analysis of $\mathrm{lb}(f)_r$ in Theorem 7.1. Indeed, as will become clear later, the condition put on ξ^n_{r+1} follows from a technical argument (see Lemma 7.13), which is not required in the proof of Theorem 7.3.

Asymptotic analysis for both hierarchies. The results above show that the relative error of both hierarchies is bounded asymptotically by the function $\varphi(t)$ from (7.5) in the regime $r \approx t \cdot n$. This is summarized in the following corollary, which can be seen as an asymptotic version of Theorem 7.1 and Theorem 7.3.

COROLLARY 7.4. *Fix $d \leq n$ and for $n, r \in \mathbb{N}$ write*

$$E_{(r)}(n) := \sup_{f \in \mathbb{R}[\mathbf{x}]_d} \big\{ f_{\min} - \mathrm{lb}(f)_r : \|f\|_\infty = 1 \big\},$$
$$E^{(r)}(n) := \sup_{f \in \mathbb{R}[\mathbf{x}]_d} \big\{ \mathrm{ub}(f)_r - f_{\min} : \|f\|_\infty = 1 \big\}.$$

Let C_d be the constant of Theorem 7.1 and let $\varphi(t)$ be the function from (7.5). Then, for any $t \in [0, 1/2]$, we have:

$$\lim_{r/n \to t} E^{(r)}(n) \leq C_d \cdot \varphi(t)$$

and, if $d(d+1) \cdot \varphi(t) \leq 1/2$, we also have:

$$\lim_{r/n \to t} E_{(r)}(n) \leq 2 \cdot C_d \cdot \varphi(t).$$

Here, the limit notation $r/n \to t$ means that the claimed convergence holds for all sequences $(n_j)_j$ and $(r_j)_j$ of integers such that $\lim_{j\to\infty} n_j = \infty$ and $\lim_{j\to\infty} r_j/n_j = t$.

We close with some remarks. First, note that $\varphi(1/2) = 0$. Hence Corollary 7.4 tells us that the relative error of both hierarchies tends to 0 as $r/n \to 1/2$. We thus 'asymptotically' recover the exactness result (7.2) of [**STKI17**].

Our results in Theorems 7.1 and 7.3 and Corollary 7.4 extend directly to the case of polynomial optimization over the discrete cube $\{\pm 1\}^n$ instead of the binary cube $\mathbb{B}^n = \{0,1\}^n$, as can easily be seen by applying a change of variables $x \in \{0,1\} \mapsto 2x - 1 \in \{\pm 1\}$. In addition, as we show in Appendix 7.5, our results extend to the case of polynomial optimization over the q-ary cube $\{0, 1, \ldots, q-1\}^n$ for $q > 2$.

Overview of the proof. As mentioned, we follow the proof strategy of Chapter 6. For convenience, we sketch here the specialization of this method to the setting of the binry cube. Let $f \in \mathbb{R}[\mathbf{x}]_d$ be the polynomial with degree d for which we wish to analyze the bounds $\mathrm{lb}(f)_r$ and $\mathrm{ub}(f)_r$. After rescaling, and up to a change of coordinates, we may assume w.l.o.g. that f attains its minimum over $\mathbb{B}^n$ at $0 \in \mathbb{B}^n$ and that $f_{\min} = 0$ and $f_{\max} = 1$. So we have $\|f\|_\infty := \max_{\mathbf{x}\in\mathbb{B}^n} |f(\mathbf{x})| = 1$. To simplify notation, we will make these assumptions throughout this chapter.

Let μ be the uniform probability measure on $\mathbb{B}^n$. A kernel $\mathrm{K} : \mathbb{B}^n \times \mathbb{B}^n \to \mathbb{R}$ induces a linear operator $\mathbf{K}$ on the space of polynomials $\mathcal{P}(\mathbb{B}^n)$ by:

$$\mathbf{K}p(\mathbf{x}) := \int_{\mathbb{B}^n} p(\mathbf{y})\mathrm{K}(\mathbf{x},\mathbf{y})d\mu(\mathbf{y}) = \frac{1}{2^n}\sum_{\mathbf{y}\in\mathbb{B}^n} p(\mathbf{y})\mathrm{K}(\mathbf{x},\mathbf{y}).$$

Recall from Chapter 6 that in order to analyze the lower bounds $\mathrm{lb}(f)_r$, it suffices to construct a kernel K whose associated operator $\mathbf{K}$ is invertible and satisfies:

$$\mathbf{K}(1) = 1, \tag{P1}$$

$$\mathbf{K}p \in \mathcal{Q}(\mathbb{B}^n)_{2r} \quad \text{for all } p \in \mathcal{P}_+(\mathbb{B}^n)_d \tag{P2}$$

$$\max_{\mathbf{x}\in\mathbb{B}^n} |\mathbf{K}^{-1}f(\mathbf{x}) - f(\mathbf{x})| \le \varepsilon. \tag{P3}$$

Lemma 7.5 (Specialization of Lemma 6.1). *Let* K *be a kernel whose associated operator satisfies* (P1)*,* (P2) *and* (P3)*. Then we have* $f_{\min} - \mathrm{lb}(f)_r \le \varepsilon$*.*

Proof. Writing $\tilde{f}(\mathbf{x}) = f(\mathbf{x}) + \varepsilon$, we have:

$$\|\mathbf{K}^{-1}\tilde{f} - \tilde{f}\|_\infty = \|\mathbf{K}^{-1}f - f\|_\infty \le \varepsilon\|f\|_\infty = \varepsilon,$$

where we use the fact that $\mathbf{K}(1) = 1 = \mathbf{K}^{-1}(1)$ for the first equality and (P3) for the inequality. We then see that $\mathbf{K}^{-1}\tilde{f}(\mathbf{x}) \ge \tilde{f}(\mathbf{x}) - \varepsilon = f(\mathbf{x}) \ge f_{\min} = 0$ on $\mathbb{B}^n$, and so (P2) implies that:

$$f(\mathbf{x}) + \varepsilon = \mathbf{K}\mathbf{K}^{-1}(f + \varepsilon) \in \mathcal{Q}(\mathbb{B}^n),$$

or in other words that $f_{\min} - \mathrm{lb}(f)_r \le \varepsilon$. □

As we see below, the key feature of this setting is that the Christoffel-Darboux kernel CD_{2r} is given by:

$$\mathrm{CD}_{2r}(\mathbf{x},\mathbf{y}) = \sum_{k=0}^{2r} \mathcal{K}_k^{(n)}(d(\mathbf{x},\mathbf{y})) \quad (\mathbf{x},\mathbf{y}\in\mathbb{B}^n),$$

where $\mathcal{K}_k^{(n)}$ is the Krawtchouk polynomial of degree k, and $d(\mathbf{x},\mathbf{y}) = \|\mathbf{x}-\mathbf{y}\|_1$ denotes the *Hamming distance* between $\mathbf{x}$ and $\mathbf{y}$. The idea is then to consider a kernel K on $\mathbb{B}^n$ of the form:

$$\mathrm{K}(\mathbf{x},\mathbf{y}) = u^2(d(\mathbf{x},\mathbf{y})),$$

where $u \in \mathbb{R}[x]_r$ is a univariate polynomial of degree at most r. This kernel clearly lies in $\mathcal{Q}(\mathbb{B}^n)_{2r}$ for fixed $\mathbf{y} \in \mathbb{B}^n$. By Lemma 6.2, this implies that its associated operator $\mathbf{K}$ satisfies (P2). Furthermore, if we write $u^2(x) =$

$\sum_{k=0}^{2r} \lambda_k \mathcal{K}_k^{(n)}(x)$ in the basis of Krawtchouk polynomials, we then have:

$$\mathrm{K}(\mathbf{x}, \mathbf{y}) = u^2(d(\mathbf{x}, \mathbf{y})) = \sum_{k=0}^{2r} \lambda_k \mathcal{K}_k^{(n)}(d(\mathbf{x}, \mathbf{y})) = \mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda),$$

where $\mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda)$ is the perturbed Christoffel-Darboux kernel (6.7). The eigenvalues of $\mathbf{K}$ are thus equal to $\lambda_0, \lambda_1, \ldots, \lambda_{2r}$. As we have seen in Chapter 6, $\mathbf{K}$ therefore satisfies (P1) if $\lambda_0 = 1$, and (P3) if $\lambda_k \approx 1$ for $1 \leq k \leq d$ (we make this precise in Section 7.2).

Interestingly, the problem of finding a polynomial u for which the coefficients λ_k satisfy these properties reduces to analyzing the quality of the measure-based upper bounds in a particular univariate setting. In order to perform this analysis and conclude the proof of Theorem 7.1, we make use of the connection between the upper bounds and roots of orthogonal polynomials (in this case the Krawtchouk polynomials) mentioned first in Chapter 2.

Organization. The rest of the chapter is structured as follows. We review the necessary background on Fourier analysis on the binary cube in Section 7.1. Then, in Section 7.2, we give a proof of Theorem 7.1. In Section 7.3, we discuss how to generalize the proofs of Section 7.2 to obtain Theorem 7.3. In Section 7.4, we give the proof of a technical lemma needed in the proof of Theorem 7.1. Finally, we we indicate in Section 7.5 how our arguments extend to the case of polynomial optimization over the q-ary hypercube $\{0, 1, \ldots, q-1\}^n$ for $q > 2$.

7.1. Preliminaries

In this section, we cover some standard Fourier analysis on the binary cube, most of which can also be found in Chapter 1. We also prove some small statements on Krawtchouk polynomials that we will need later.

7.1.1. Notations.

For $n \in \mathbb{N}$, we write $\mathbb{B}^n = \{0,1\}^n$ for the binary hypercube of dimension n. We let μ denote the uniform probability measure on $\mathbb{B}^n$, given by $\mu = \frac{1}{2^n} \sum_{\mathbf{x} \in \mathbb{B}^n} \delta_{\mathbf{x}}$, where $\delta_{\mathbf{x}}$ is the Dirac measure at $\mathbf{x}$. Further, we write $|\mathbf{x}| = \sum_i \mathbf{x}_i = |\{i \in [n] : \mathbf{x}_i = 1\}|$ for the *Hamming weight* of $\mathbf{x} \in \mathbb{B}^n$, and $d(\mathbf{x}, \mathbf{y}) = |\{i \in [n] : \mathbf{x}_i \neq \mathbf{y}_i\}|$ for the *Hamming distance* between $\mathbf{x}, \mathbf{y} \in \mathbb{B}^n$. We let $\mathrm{Sym}(n)$ denote the set of permutations of the set $[n] = \{1, \ldots, n\}$.

We consider polynomials $p : \mathbb{B}^n \to \mathbb{R}$ on $\mathbb{B}^n$. The space $\mathcal{P}(\mathbb{B}^n)$ of such polynomials is given by the quotient ring of $\mathbb{R}[\mathbf{x}]$ over the equivalence relation $p \sim q$ if $p(\mathbf{x}) = q(\mathbf{x})$ on $\mathbb{B}^n$. In other words, $\mathcal{P}(\mathbb{B}^n) = \mathbb{R}[\mathbf{x}]/\mathcal{I}$, where $\mathcal{I}$ is the ideal generated by the polynomials $\mathbf{x}_i - \mathbf{x}_i^2$ for $i \in [n]$, which can also be seen as the vector space spanned by the (multilinear) polynomials $\prod_{i \in I} \mathbf{x}_i$ for $I \subseteq [n]$.

For $a \leq b \in \mathbb{N}$, we let $[a : b]$ denote the set of integers $a, a+1, \ldots, b$.

7.1.2. The character basis. Let $\langle \cdot, \cdot \rangle_\mu$ be the inner product on $\mathcal{P}(\mathbb{B}^n)$ given by:

$$\langle p, q \rangle_\mu = \int_{\mathbb{B}^n} p(\mathbf{x})q(\mathbf{x})d\mu(\mathbf{x}) = \frac{1}{2^n} \sum_{\mathbf{x} \in \mathbb{B}^n} p(\mathbf{x})q(\mathbf{x}).$$

The space $\mathcal{P}(\mathbb{B}^n)$ has an orthonormal basis w.r.t. $\langle \cdot, \cdot \rangle_\mu$ given by the *characters*:

$$\chi_a(\mathbf{x}) := (-1)^{\mathbf{x} \cdot a} = \prod_{i : a_i = 1} (1 - 2\mathbf{x}_i) \quad (a \in \mathbb{B}^n). \tag{7.7}$$

In other words, the set $\{\chi_a : a \in \mathbb{B}^n\}$ of all characters on $\mathbb{B}^n$ forms a basis for $\mathcal{P}(\mathbb{B}^n)$ and

$$\langle \chi_a, \chi_b \rangle_\mu = \frac{1}{2^n} \sum_{\mathbf{x} \in \mathbb{B}^n} \chi_a(\mathbf{x})\chi_b(\mathbf{x}) = \delta_{a,b} \quad \forall a, b \in \mathbb{B}^n. \tag{7.8}$$

Then any polynomial $p \in \mathcal{P}(\mathbb{B}^n)$ can be expressed in the basis of characters, known as its *Fourier expansion*:

$$p(\mathbf{x}) = \sum_{a \in \mathbb{B}^n} \widehat{p}(a)\chi_a(\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{B}^n \tag{7.9}$$

with *Fourier coefficients* $\widehat{p}(a) := \langle p, \chi_a \rangle_\mu \in \mathbb{R}$.

The group $\mathrm{Aut}(\mathbb{B}^n)$ of automorphisms of $\mathbb{B}^n$ is generated by the coordinate permutations, of the form $\mathbf{x} \mapsto \sigma(\mathbf{x}) := (\mathbf{x}_{\sigma(1)}, \dots, \mathbf{x}_{\sigma(n)})$ for $\sigma \in \mathrm{Sym}(n)$, and the automorphisms corresponding to bit-flips, of the form $\mathbf{x} \in \mathbb{B}^n \mapsto \mathbf{x} \oplus a \in \mathbb{B}^n$ for $a \in \mathbb{B}^n$. If we set

$$H_k := \mathrm{span}\{\chi_a : |a| = k\} \quad (0 \leq k \leq n),$$

then each H_k is an irreducible, $\mathrm{Aut}(\mathbb{B}^n)$-invariant subspace of $\mathcal{P}(\mathbb{B}^n)$ of dimension $\binom{n}{k}$. Using (7.9), we may then decompose $\mathcal{P}(\mathbb{B}^n)$ as the direct sum:

$$\mathcal{P}(\mathbb{B}^n) = H_0 \perp H_1 \perp \cdots \perp H_n,$$

where the subspaces H_k are pairwise orthogonal w.r.t. $\langle \cdot, \cdot \rangle_\mu$. In fact, we have that $\mathcal{P}(\mathbb{B}^n)_d = H_0 \perp H_1 \perp \cdots \perp H_d$ for all $d \leq n$, and we may thus write any $p \in \mathcal{P}(\mathbb{B}^n)_d$ (in a unique way) as

$$p = p_0 + p_1 + \cdots + p_d \quad (p_k \in H_k).$$

The polynomials $p_k \in H_k$ $(k = 0, \dots, d)$ are known as the *harmonic components* of p and the decomposition (7.1.2) as the *harmonic decomposition* of p. We will make extensive use of this decomposition throughout.

Let $\mathrm{St}(0) \subseteq \mathrm{Aut}(\mathbb{B}^n)$ be the set of automorphisms fixing $0 \in \mathbb{B}^n$, which consists of the coordinate permutations $\mathbf{x} \mapsto \sigma(\mathbf{x}) = (\mathbf{x}_{\sigma(1)}, \dots, \mathbf{x}_{\sigma(n)})$ for $\sigma \in \mathrm{Sym}(n)$. The subspace of functions in H_k that are invariant under $\mathrm{St}(0)$ is one-dimensional and it is spanned by the function

$$X_k(\mathbf{x}) := \sum_{|a|=k} \chi_a(\mathbf{x}). \tag{7.10}$$

These functions X_k are known as the *zonal spherical functions* with pole $0 \in \mathbb{B}^n$.

7.1.3. Krawtchouk polynomials. For $k \in \mathbb{N}$, the *Krawtchouk polynomial* of degree k (and with parameter n) is the univariate polynomial in x given by:

$$\mathcal{K}_k^{(n)}(x) := \sum_{i=0}^{k} (-1)^i \binom{x}{i} \binom{n-x}{k-i}. \tag{7.11}$$

The Krawtchouk polynomials form an orthogonal basis for $\mathbb{R}[x]$ with respect to the inner product given by the following discrete probability measure on the set $[0:n] = \{0, 1, \ldots, n\}$:

$$\omega := \frac{1}{2^n} \sum_{t=0}^{n} w(x)\delta_x, \text{ with } w(x) := \binom{n}{x}.$$

Indeed, for all $k, k' \in \mathbb{N}$ we have:

$$\langle \mathcal{K}_k^{(n)}, \mathcal{K}_{k'}^{(n)} \rangle_\omega := \frac{1}{2^n} \sum_{x=0}^{n} \mathcal{K}_k^{(n)}(x) \mathcal{K}_{k'}^{(n)}(x) w(x) = \delta_{k,k'} \binom{n}{k}. \tag{7.12}$$

The following (well-known) lemma explains the connection between the Krawtchouk polynomials and the character basis on $\mathcal{P}(\mathbb{B}^n)$.

Lemma 7.6. *Let $x \in [0:n]$ and choose $\mathbf{x}, \mathbf{y} \in \mathbb{B}^n$ so that $d(\mathbf{x}, \mathbf{y}) = x$. Then for any $0 \leq k \leq n$ we have:*

$$\mathcal{K}_k^{(n)}(x) = \sum_{|a|=k} \chi_a(\mathbf{x})\chi_a(\mathbf{y}). \tag{7.13}$$

In particular, we have:

$$\mathcal{K}_k^{(n)}(x) = \sum_{|a|=k} \chi_a(1^x 0^{n-x}) = X_k(1^x 0^{n-x}), \tag{7.14}$$

where $1^x 0^{n-x} \in \mathbb{B}^n$ is given by $(1^x 0^{n-x})_i = 1$ if $1 \leq i \leq x$ and $(1^x 0^{n-x})_i = 0$ if $x + 1 \leq i \leq n$.

Proof. Noting that $\chi_a(\mathbf{x})\chi_a(\mathbf{y}) = \chi_a(\mathbf{x}+\mathbf{y})$ and $|\mathbf{x}+\mathbf{y}| = d(\mathbf{x}, \mathbf{y}) = x$, we have:

$$\begin{aligned}\sum_{|a|=k} \chi_a(\mathbf{x})\chi_a(\mathbf{y}) &= \sum_{i=0}^{k} (-1)^i \cdot \#\{|a| = k : a \cdot (\mathbf{x}+\mathbf{y}) = i\} \\ &= \sum_{i=0}^{k} (-1)^i \binom{x}{i} \binom{n-x}{k-i} = \mathcal{K}_k^{(n)}(x).\end{aligned}$$

□

From this, we see that any polynomial $p \in \mathcal{P}(\mathbb{B}^n)_d$ that is invariant under the action of $\mathrm{St}(0)$ is of the form $\sum_{k=1}^d \lambda_k \mathcal{K}_k^{(n)}(|\mathbf{x}|)$ for some scalars λ_k, and thus $p(\mathbf{x}) = u(|\mathbf{x}|)$ for some univariate polynomial $u \in \mathbb{R}[x]_d$.

It will sometimes be convenient to work with a different normalization of the Krawtchouk polynomials, given by:

$$\widehat{\mathcal{K}}_k^n(x) := \mathcal{K}_k^{(n)}(x)/\mathcal{K}_k^{(n)}(0) \quad (k \in \mathbb{N}). \tag{7.15}$$

So $\widehat{\mathcal{K}}_k^n(0) = 1$. Note that for any $k \in \mathbb{N}$, we have

$$\|\mathcal{K}_k^{(n)}\|_\omega^2 := \langle \mathcal{K}_k^{(n)}, \mathcal{K}_k^{(n)} \rangle_\omega = \binom{n}{k} = \mathcal{K}_k^{(n)}(0),$$

meaning that $\widehat{\mathcal{K}}_k^n(x) = \mathcal{K}_k^{(n)}(x)/\|\mathcal{K}_k^{(n)}\|_\omega^2$.

Finally we give a short proof of two basic facts on Krawtchouk polynomials that will be used below.

LEMMA 7.7. *We have:*

$$\widehat{\mathcal{K}}_k^n(x) \leq \widehat{\mathcal{K}}_0^n(x) = 1$$

for all $0 \leq k \leq n$ *and* $x \in [0:n]$.

PROOF. Given $x \in [0:n]$ consider an element $\mathbf{x} \in \mathbb{B}^n$ with Hamming weight x, for instance the element $1^x 0^{n-x}$ from Lemma 7.6. By (7.14) we have

$$\mathcal{K}_k^{(n)}(x) = \sum_{|a|=k} \chi_a(\mathbf{x}) \leq \binom{n}{k} = \mathcal{K}_k^{(n)}(0),$$

making use of the fact that $|\chi_a(\mathbf{x})| = 1$ for all $a \in \mathbb{B}^n$. □

LEMMA 7.8. *We have:*

$$\begin{aligned} |\widehat{\mathcal{K}}_k^n(x) - \widehat{\mathcal{K}}_k^n(x+1)| &\leq \frac{2k}{n}, \qquad && (x = 0, 1, \ldots, n-1) \\ |\widehat{\mathcal{K}}_k^n(x) - 1| &\leq \frac{2k}{n} \cdot x \qquad && (x = 0, 1, \ldots, n) \end{aligned} \tag{7.16}$$

for all $0 \leq k \leq n$.

PROOF. Let $x \in [0:n-1]$ and $0 < k \leq d$. Consider the elements $1^x 0^{n-x}$ and $1^{x+1} 0^{n-x-1}$ of $\mathbb{B}^n$ from Lemma 7.6. We have:

$$\begin{aligned} |\mathcal{K}_k^{(n)}(x) - \mathcal{K}_k^{(n)}(x+1)| &\overset{(7.14)}{=} |\sum_{|a|=k} \chi_a(1^x 0^{n-x}) - \chi_a(1^{x+1} 0^{n-x-1})| \\ &\leq 2 \cdot \#\{a \in \mathbb{B}^n : |a| = k, a_{x+1} = 1\} = 2\binom{n-1}{k-1}, \end{aligned}$$

where for the inequality we note that $\chi_a(1^x 0^{n-x}) = \chi_a(1^{x+1} 0^{n-x-1})$ if $a_{x+1} = 0$. As $\mathcal{K}_k^{(n)}(0) = \binom{n}{k}$, this implies that:

$$|\widehat{\mathcal{K}}_k^n(x) - \widehat{\mathcal{K}}_k^n(x+1)| \leq 2\binom{n-1}{k-1} / \binom{n}{k} = \frac{2k}{n}.$$

This shows the first inequality of (7.16). The second inequality follows using the triangle inequality, a telescope summation argument and the fact that $\widehat{\mathcal{K}}_k^n(0) = 1$. □

7.1.4. Invariant kernels and the Funk-Hecke formula. Given a univariate polynomial $u \in \mathbb{R}[x]$ of degree r with $2r \geq d$, consider the kernel $\mathrm{K} : \mathbb{B}^n \times \mathbb{B}^n \to \mathbb{R}$ defined by

$$\mathrm{K}(\mathbf{x}, \mathbf{y}) := u^2(d(\mathbf{x}, \mathbf{y})). \tag{7.17}$$

A kernel of the form (7.17) coincides with a polynomial of degree $2\deg(u)$ in $\mathbf{x}$ on the binary cube $\mathbb{B}^n$, as $d(\mathbf{x}, \mathbf{y}) = \sum_i (\mathbf{x}_i + \mathbf{y}_i - 2\mathbf{x}_i\mathbf{y}_i)$ for $\mathbf{x}, \mathbf{y} \in \mathbb{B}^n$. Furthermore, it is invariant under $\mathrm{Aut}(\mathbb{B}^n)$, in the sense that:

$$\mathrm{K}(\mathbf{x}, \mathbf{y}) = \mathrm{K}(\pi(\mathbf{x}), \pi(\mathbf{y})) \quad \forall \mathbf{x}, \mathbf{y} \in \mathbb{B}^n, \pi \in \mathrm{Aut}(\mathbb{B}^n).$$

The kernel K acts as a linear operator $\mathbf{K} : \mathcal{P}(\mathbb{B}^n) \to \mathcal{P}(\mathbb{B}^n)$ by:

$$\mathbf{K}p(\mathbf{x}) := \int_{\mathbb{B}^n} p(\mathbf{y})\mathrm{K}(\mathbf{x}, \mathbf{y})d\mu(\mathbf{y}) = \frac{1}{2^n} \sum_{\mathbf{y} \in \mathbb{B}^n} p(\mathbf{y})\mathrm{K}(\mathbf{x}, \mathbf{y}).$$

We may expand the polynomial $u^2 \in \mathbb{R}[x]_{2r}$ in the basis of Krawtchouk polynomials as:

$$u^2(x) = \sum_{k=0}^{2r} \lambda_k \mathcal{K}_k^{(n)}(x) \quad (\lambda_k \in \mathbb{R}). \tag{7.18}$$

As we show now, the eigenvalues of the operator $\mathbf{K}$ are given precisely by the coefficients λ_k occurring in this expansion. This relation is analogous to the classical Funk-Hecke formula for spherical harmonics, see also Chapter 1.

THEOREM 7.9 (Funk-Hecke). *Let $p \in \mathcal{P}(\mathbb{B}^n)_d$ with harmonic decomposition $p = p_0 + p_1 + \cdots + p_d$. Then we have:*

$$\mathbf{K}p = \lambda_0 p_0 + \lambda_1 p_1 + \cdots + \lambda_d p_d. \tag{7.19}$$

PROOF. It suffices to show that $\mathbf{K}\chi_z = \lambda_{|z|}\chi_z$ for all $z \in \mathbb{B}^n$. So we compute for $\mathbf{x} \in \mathbb{B}^n$:

$$\begin{aligned}
\mathbf{K}\chi_z(\mathbf{x}) = \frac{1}{2^n}\sum_{\mathbf{y}\in\mathbb{B}^n}\chi_z(\mathbf{y})u^2(d(\mathbf{x},\mathbf{y})) &\overset{(7.18)}{=} \frac{1}{2^n}\sum_{\mathbf{y}\in\mathbb{B}^n}\chi_z(\mathbf{y})\sum_{i=0}^{2r}\lambda_i\mathcal{K}_i^{(n)}(d(\mathbf{x},\mathbf{y})) \\
&\overset{(7.13)}{=} \sum_{i=0}^{2r}\lambda_i\sum_{\mathbf{y}\in\mathbb{B}^n}\chi_z(\mathbf{y})\sum_{|a|=i}\chi_a(\mathbf{x})\chi_a(\mathbf{y}) \\
&= \sum_{i=0}^{2r}\lambda_i\sum_{|a|=i}\Big(\sum_{\mathbf{y}\in\mathbb{B}^n}\chi_z(\mathbf{y})\chi_a(\mathbf{y})\Big)\chi_a(\mathbf{x}) \\
&\overset{(7.8)}{=} \frac{1}{2^n}\sum_{i=0}^{2r}\lambda_i\sum_{|a|=i}2^n\delta_{z,a}\chi_a(\mathbf{x}) \\
&= \lambda_{|z|}\chi_z(\mathbf{x}).
\end{aligned}$$

□

Finally, we note that since the Krawtchouk polynomials form an *orthogonal* basis for $\mathbb{R}[x]$, we may express the coefficients λ_k in the decomposition (7.18) of u^2 in the following way:

$$\lambda_k = \langle\mathcal{K}_k^{(n)}, u^2\rangle_\omega \;/\; \|\mathcal{K}_k^{(n)}\|_\omega^2 = \langle\widehat{\mathcal{K}}_k^n, u^2\rangle_\omega. \tag{7.20}$$

In addition, since in view of Lemma 7.7 we have $\widehat{\mathcal{K}}_k^n(x) \le \widehat{\mathcal{K}}_0^n(x)$ for all $x \in [0:n]$, it folllows that

$$\lambda_k \le \lambda_0 \quad \text{for } 0 \le k \le 2r. \tag{7.21}$$

7.2. Proof of main result

We consider a polynomial f of degree $d \in \mathbb{N}$ to be minimized over $\mathbb{B}^n$, which we assume w.l.o.g. to satisfy $f_{\min} = f(0) = 0$ and $\|f\|_\infty = 1$. Let $u \in \mathbb{R}[x]$ be a univariate polynomial of degree r (which we select later) with $2r \ge d$. Recall that we decompose u^2 in the basis of Krawtchouck polynomials as:

$$u^2(x) = \sum_{k=0}^{2r}\lambda_k\mathcal{K}_k^{(n)}(x) \quad (\lambda_k \in \mathbb{R}). \tag{7.22}$$

Consider the kernel $\mathrm{K}(\mathbf{x},\mathbf{y}) = u^2(d(\mathbf{x},\mathbf{y}))$ and its associated linear operator $\mathbf{K}$, which has eigenvalues $\lambda_0, \lambda_1, \ldots, \lambda_{2r}$ by the Funk-Hecke formula (7.19). In order to prove Theorem 7.1, it suffices to show that u may be chosen so that $\mathbf{K}$ satisfies (P1), (P2) and (P3).

Let us first note that the polynomial $\mathbf{x} \to \mathrm{K}(\mathbf{x},\mathbf{y})$ lies in $\mathcal{Q}(\mathbb{B}^n)_{2r}$ for each $\mathbf{y} \in \mathbb{B}^n$ by definition. Therefore, $\mathbf{K}$ always satisfies (P2).

LEMMA 7.10 (Specialization of Lemma 6.2). *Let* K *be a polynomial kernel on* $\mathbb{B}^n$ *and assume that* $\mathrm{K}(\mathbf{x}, \mathbf{y})$ *lies in* $\mathcal{Q}(\mathbb{B}^n)_{2r}$ *for each* $\mathbf{y} \in \mathbb{B}^n$. *Then the associated operator* $\mathbf{K}$ *satisfies* (P2).

PROOF. Let $p \in \mathbb{R}[\mathbf{x}]$ be a polynomial, and assume that $p(\mathbf{x}) \geq 0$ for all $\mathbf{x} \in \mathbb{B}^n$. Then we have:

$$\mathbf{K}p(\mathbf{x}) := \frac{1}{2^n} \sum_{\mathbf{y} \in \mathbb{B}^n} \underbrace{\mathrm{K}(\mathbf{x}, \mathbf{y})}_{\in \mathcal{Q}(\mathbb{B}^n)} \underbrace{p(\mathbf{y})}_{\geq 0},$$

which lies in the cone $\mathcal{Q}(\mathbb{B}^n)_{2r}$. □

Next, we note that $\mathbf{K}(1) = \lambda_0$, and so (P1) is satisfied precisely when $\lambda_0 = 1$. It remains to consider (P3). Recall that we are interested there in bounding the quantity:

$$\max_{\mathbf{x} \in \mathbb{B}^n} |\mathbf{K}^{-1} f(\mathbf{x}) - f(\mathbf{x})|.$$

Our approach consists of two parts. First, we relate this quantity to the coefficients λ_k in the decomposition (7.22). Then, using this relation and the connection between Lasserre's *upper* bounds and extremal roots of orthogonal polynomials outlined in Section 2.2, we show that u may be chosen such that the quantity is of the order ξ_{r+1}^n / n, where ξ_{r+1}^n is the smallest root of the degree $r + 1$ Krawtchouk polynomial (with parameter n).

7.2.1. Expressing (P3) in terms of the coefficients λ_k. We need the following technical lemma, which bounds the sup-norm $\|p_k\|_\infty$ of the harmonic components p_k of a polynomial $p \in \mathcal{P}(\mathbb{B}^n)$ in terms of $\|p\|_\infty$, the sup-norm of p itself. The key point is that this bound is independent of the dimension n. We delay the proof which is rather technical to Section 7.4.

LEMMA 7.11. *There exists a constant* $\gamma_d > 0$, *depending only on* d, *such that for any* $p = p_0 + p_1 + \ldots + p_d \in \mathcal{P}(\mathbb{B}^n)_d$, *we have:*

$$\|p_k\|_\infty \leq \gamma_d \|p\|_\infty \text{ for all } 0 \leq k \leq d.$$

COROLLARY 7.12. *Assume that* $\lambda_0 = 1$ *and* $\lambda_k \neq 0$ *for* $1 \leq k \leq d$. *Then we have:*

$$\max_{\mathbf{x} \in \mathbb{B}^n} |\mathbf{K}^{-1} f(\mathbf{x}) - f(\mathbf{x})| \leq \gamma_d \cdot \Lambda, \text{ where } \Lambda := \sum_{k=1}^{d} |\lambda_k^{-1} - 1|. \tag{7.23}$$

PROOF. By assumption, the operator $\mathbf{K}$ is invertible and, in view of Funk-Hecke formula (7.19), its inverse is given by: $\mathbf{K}^{-1} p = \sum_{i=0}^{d} \lambda_i^{-1} p_k$ for any

$p = p_0 + p_1 + \ldots + p_d \in \mathcal{P}(\mathbb{B}^n)_d$. For each $\mathbf{x} \in \mathbb{B}^n$, we therefore have:

$$\begin{aligned}|\mathbf{K}^{-1}f(\mathbf{x}) - f(\mathbf{x})| = |\sum_{k=1}^{d}(\lambda_i^{-1} - 1)f_k(\mathbf{x})| &\leq \sum_{k=1}^{d}|\lambda_k^{-1} - 1|\|f_k\|_\infty \\ &\leq \sum_{k=1}^{d}|\lambda_k^{-1} - 1| \cdot \gamma_d,\end{aligned}$$

where we use Lemma 7.11 for the last inequality. □

The expression Λ in (7.23) is difficult to analyze. Therefore, following [**FF21**], we consider instead the simpler expression:

$$\tilde{\Lambda} := \sum_{k=1}^{d}(1 - \lambda_k) = d - \sum_{k=1}^{d}\lambda_k,$$

which is linear in the λ_k. Under the assumption that $\lambda_0 = 1$, we have $\lambda_k \leq \lambda_0 = 1$ for all k (recall relation (7.21)). Thus, Λ and $\tilde{\Lambda}$ are both minimized when the λ_k are close to 1. The following lemma makes this precise.

LEMMA 7.13. *Assume that* $\lambda_0 = 1$ *and that* $\tilde{\Lambda} \leq 1/2$. *Then we have* $\Lambda \leq 2\tilde{\Lambda}$, *and thus that:*

$$\max_{\mathbf{x} \in \mathbb{B}^n} |\mathbf{K}^{-1}f(\mathbf{x}) - f(\mathbf{x})| \leq 2\gamma_d \cdot \tilde{\Lambda}.$$

PROOF. As we assume $\tilde{\Lambda} \leq 1/2$, we must have $1/2 \leq \lambda_k \leq 1$ for all k. Therefore, we may write:

$$\Lambda = \sum_{k=1}^{d}|\lambda_k^{-1} - 1| = \sum_{k=1}^{d}|(1 - \lambda_k)/\lambda_k| = \sum_{k=1}^{d}(1 - \lambda_k)/\lambda_k \leq 2\sum_{k=1}^{d}(1 - \lambda_k) = 2\tilde{\Lambda}.$$

□

7.2.2. Optimizing the choice of the univariate polynomial u. In light of Lemma 7.13, and recalling (7.20), we wish to find a univariate polynomial $u \in \mathbb{R}[x]_r$ for which:

$$\lambda_0 = \langle 1, u^2 \rangle_\omega = 1, \text{ and}$$

$$\tilde{\Lambda} = d - \sum_{k=1}^{d}\lambda_k = d - \sum_{k=1}^{d}\langle \widehat{\mathcal{K}}_k^n, u^2 \rangle_\omega \text{ is small.}$$

Unpacking the definition of $\langle \cdot, \cdot \rangle_\omega$, we thus need to solve the following optimization problem:

$$\begin{aligned}&\inf_{u \in \mathbb{R}[x]_r} \left\{ \int g(x) \cdot u^2(x) d\omega(x) : \int u^2(x) d\omega(x) = 1 \right\}, \\ &\text{where } g(x) := d - \sum_{k=1}^{d} \widehat{\mathcal{K}}_k^n(x).\end{aligned} \tag{7.24}$$

(Indeed $\int gu^2 d\omega = \langle g, u^2\rangle_\omega = \tilde{\Lambda}$ and $\int u^2 d\omega = \langle 1, u^2\rangle_\omega$.) We recognize that this is exactly[1] the program that defines the *upper* bound $\mathrm{ub}(g, [0:n], \omega)_r$ for the minimization of g on $[0:n]$! Hence the optimal value of (7.24) is equal to $\mathrm{ub}(g)_r := \mathrm{ub}(g, [0:n], \omega)_r$ and, using Lemma 7.13, we may conclude the following.

THEOREM 7.14. *Let g be as in (7.24). Assume that $\mathrm{ub}(g)_r \leq 1/2$. Then there exists a polynomial $u \in \mathbb{R}[x]_r$ such that $\lambda_0 = 1$ and:*

$$\max_{\mathbf{x}\in\mathbb{B}^n} |\mathbf{K}^{-1} f(\mathbf{x}) - f(\mathbf{x})| \leq 2\gamma_d \cdot \mathrm{ub}(g)_r.$$

It remains, then, to analyze $\mathrm{ub}(g)_r$. For this purpose, we follow a technique outlined in Section 2.3. Note that $g_{\min} = g(0) = 0$. We first show that g can be upper bounded by its linear Taylor approximation at $x = 0$.

LEMMA 7.15. *We have:*

$$g(x) \leq \widehat{g}(x) := d(d+1) \cdot (x/n) \quad \forall x \in [0:n].$$

Furthermore, the minimum $\hat{g}_{\min}$ of $\hat{g}$ on $[0:n]$ clearly satisfies $\widehat{g}_{\min} = \widehat{g}(0) = g(0) = g_{\min}$.

PROOF. Using (7.16), we find for each $k \leq n$ that:

$$\widehat{\mathcal{K}}_k^n(x) \geq \widehat{\mathcal{K}}_k^n(0) - \frac{2k}{n} \cdot x = 1 - \frac{2k}{n} \cdot x \quad \forall x \in [0:n].$$

Therefore, we have:

$$g(x) := d - \sum_{k=1}^{d} \widehat{\mathcal{K}}_k^n(x) \leq \sum_{k=1}^{d} \frac{2k}{n} \cdot x = d(d+1) \cdot (x/n) \quad \forall x \in [0:n].$$

□

LEMMA 7.16. *We have:*

$$\mathrm{ub}(g)_r \leq d(d+1) \cdot (\xi_{r+1}^n / n),$$

where ξ_{r+1}^n is the smallest root of the Krawtchouk polynomial $\mathcal{K}_{r+1}^{(n)}$.

PROOF. This follows immediately from Proposition 1.2, noting that the Krawtchouk polynomials are indeed orthogonal w.r.t. the measure ω on $[0:n]$ (cf. (7.12)). □

Putting things together, we may prove our main result, Theorem 7.1.

[1]Technically, the density should be allowed to be a *sum of squares*, whereas the program (7.24) requires it to be an actual square. This is no true restriction, though, since, as a straightforward convexity argument shows, the optimum solution to (2.6) can in fact always be chosen to be a square [**Las11**].

Proof of Theorem 7.1. Assume that r is large enough so that $d(d+1)\cdot(\xi_{r+1}^n/n) \le 1/2$. By Lemma 7.16, we then have

$$\mathrm{ub}(g)_r \le d(d+1)\cdot(\xi_{r+1}^n/n) \le 1/2.$$

We are thus able to choose a polynomial $u \in \mathbb{R}[x]_r$ whose associated operator $\mathbf{K}$ satisfies $\mathbf{K}(1) = 1$ and

$$\max_{\mathbf{x}\in\mathbb{B}^n} |\mathbf{K}^{-1}f(\mathbf{x}) - f(\mathbf{x})| \le 2\gamma_d \cdot d(d+1)\cdot(\xi_{r+1}^n/n).$$

That is, we may construct an operator $\mathbf{K}$ satisfying (P1), (P2) and (P3) with $\varepsilon = 2\gamma_d \cdot d(d+1)\cdot(\xi_{r+1}^n/n)$. We may use Lemma 7.5 to obtain Theorem 7.1 with constant $C_d := \gamma_d \cdot d(d+1)$.

□

7.3. The upper bounds

We turn now to analyzing the hierarchy $\mathrm{ub}(f)_r$ of upper bounds defined in (7.3) for a polynomial $f \in \mathbb{R}[\mathbf{x}]_d$ on the binary cube, whose definition is repeated for convenience:

$$\mathrm{ub}(f)_r := \inf_{s\in\Sigma[\mathbf{x}]_r} \left\{ \int_{\mathbb{B}^n} f(\mathbf{x})\cdot s(\mathbf{x})d\mu : \int_{\mathbb{B}^n} s(\mathbf{x})d\mu = 1 \right\} \ge f_{\min}.$$

In principle, our analysis follows immediately from the proof of Theorem 7.1; see Section 6.2. For exposition, we provide here a more direct argument.

As before, we may assume w.l.o.g. that $f_{\min} = f(0) = 0$ and that $f_{\max} = 1$. To facilitate the analysis of the bounds $\mathrm{ub}(f)_r$, the idea is to restrict in (7.3) to polynomials $s(\mathbf{x})$ that are invariant under the action of $\mathrm{St}(0) \subseteq \mathrm{Aut}(\mathbb{B}^n)$, i.e., depending only on the Hamming weight $|\mathbf{x}|$. Such polynomials are of the form $s(\mathbf{x}) = u(|\mathbf{x}|)$ for some univariate polynomial $u \in \mathbb{R}[x]$. Hence this leads to the following, weaker hierarchy, where we now optimize over *univariate* sums-of-squares:

$$\mathrm{ub}(f)_r^{\mathrm{sym}} := \inf_{u\in\Sigma[x]_r} \left\{ \int_{\mathbb{B}^n} f(\mathbf{x})\cdot u(|\mathbf{x}|)d\mu(\mathbf{x}) : \int_{\mathbb{B}^n} u(|\mathbf{x}|)d\mu(\mathbf{x}) = 1 \right\}.$$

By definition, we must have $\mathrm{ub}(f)_r^{\mathrm{sym}} \ge \mathrm{ub}(f)_r \ge f_{\min}$, and so an analysis of $\mathrm{ub}(f)_r^{\mathrm{sym}}$ extends immediately to $\mathrm{ub}(f)_r$.

The main advantage of working with the hierarchy $\mathrm{ub}(f)_r^{\mathrm{sym}}$ is that we may now assume that f is itself invariant under $\mathrm{St}(0)$, after replacing f by its symmetrization:

$$\frac{1}{|\mathrm{St}(0)|}\sum_{\sigma\in\mathrm{St}(0)} f(\sigma(\mathbf{x})).$$

Indeed, for any $u \in \Sigma[x]_{2r}$, we have that:

$$\begin{aligned}\int_{\mathbb{B}^n} f(\mathbf{x})u(|\mathbf{x}|)d\mu(\mathbf{x}) &= \frac{1}{|\mathrm{St}(0)|}\sum_{\sigma\in \mathrm{St}(0)} \int_{\mathbb{B}^n} f(\sigma(\mathbf{x}))u(|\sigma(\mathbf{x})|)d\mu(\sigma(\mathbf{x})) \\ &= \int_{\mathbb{B}^n} \frac{1}{|\mathrm{St}(0)|}\sum_{\sigma\in \mathrm{St}(0)} f(\sigma(\mathbf{x}))u(|\mathbf{x}|)d\mu(\mathbf{x}).\end{aligned}$$

So we now assume that f is $\mathrm{St}(0)$-invariant, and thus we may write:

$$f(\mathbf{x}) = F(|\mathbf{x}|) \text{ for some polynomial } F(x) \in \mathbb{R}[x]_d.$$

By the definitions of the measures μ on $\mathbb{B}^n$ and ω on $[0 : n]$ we have the identities:

$$\int_{\mathbb{B}^n} u(|\mathbf{x}|)d\mu(\mathbf{x}) = \int_{[0:n]} u(x)d\omega(x),$$

$$\int_{\mathbb{B}^n} F(|\mathbf{x}|)u(|\mathbf{x}|)d\mu(\mathbf{x}) = \int_{[0:n]} F(x)u(x)d\omega(x).$$

Hence we get

$$\begin{aligned}\mathrm{ub}(f)_r^{\mathrm{sym}} &= \inf_{u\in\Sigma[x]_{2r}} \left\{ \int_{[0:n]} F(x)\cdot u(x)d\omega(x) : \int_{[0:n]} u(x)d\omega(x) = 1\right\} \\ &= \mathrm{ub}(F, [0 : n], \omega)_r.\end{aligned}$$

In other words, the behaviour of the *symmetrized* hierarchy $\mathrm{ub}(f)_r^{\mathrm{sym}}$ over the binary cube w.r.t. the uniform measure μ is captured by the behaviour of the *univariate* hierarchy $\mathrm{ub}(F, [0 : n], \omega)_r$ over $[0 : n]$ w.r.t. the discrete measure ω.

Now, we are in a position to make use again of the technique we employed at the end of Section 7.2. First we find a linear upper estimator $\widehat{F}$ for F on $[0 : n]$.

LEMMA 7.17. *We have*

$$F(x) \leq \widehat{F}(x) := d(d + 1)\cdot\gamma_d\cdot x/n \quad \forall x \in [0 : n],$$

where γ_d is the same constant as in Lemma 7.11.

PROOF. Write $F(x) = \sum_{k=0}^{d} \lambda_k \widehat{\mathcal{K}}_k^n(x)$ for some scalars λ_k. By assumption, $F(0) = 0$ and thus $\sum_{k=0}^{d} \lambda_k = 0$. We now use an analogous argument as for Lemma 7.15:

$$\begin{aligned}F(x) = \sum_{k=0}^{d} \lambda_k(\widehat{\mathcal{K}}_k^n(x) - 1) \leq \sum_{k=0}^{d} |\lambda_k||\widehat{\mathcal{K}}_k^n(x) - 1| &\overset{(7.16)}{\leq} \max_k |\lambda_k|\cdot x\cdot \sum_{k=0}^{d}\frac{2k}{n} \\ &\leq \max_k |\lambda_k|\cdot x\cdot \frac{d(d+1)}{n}.\end{aligned}$$

As $\|f\|_\infty = 1$, using Lemma 7.11, we can conclude that:

$$|\lambda_k| = \max_{x\in[0:n]} |\lambda_k\widehat{\mathcal{K}}_k^n(x)| \leq \gamma_d$$

which gives the desired result. □

In light of Lemma 7.16, we may now conclude that:

$$\mathrm{ub}(F, [0:n], \omega)_r \leq d(d+1)\gamma_d \cdot \xi_{r+1}^n / n.$$

As $\mathrm{ub}(f)_r \leq \mathrm{ub}(f)_r^{\mathrm{sym}} = \mathrm{ub}(F, [0:n], \omega)_r$, we have thus shown Theorem 7.3 with constant $C_d = d(d+1)\gamma_d$. Note that in comparison to Lemma 7.16, we only have the additional constant factor γ_d.

7.3.1. Exactness of the inner hierarchy. As is the case for the outer hierarchy, the inner hierarchy is exact when r is large enough. Whereas the outer hierarchy, however, is exact for $r \geq (n+d-1)/2$, the inner hierarchy is exact in general if and only if $r \geq n$. We give a short proof of this fact below, for reference.

LEMMA 7.18. *Let f be a polynomial on $\mathbb{B}^n$. Then $\mathrm{ub}(f)_r = f_{\min}$ for all $r \geq n$.*

PROOF. We may assume w.l.o.g. that $f(0) = f_{\min}$. Consider the interpolation polynomial:

$$s(\mathbf{x}) := \sqrt{2^n} \prod_{i=1}^{n} (1 - \mathbf{x}_i) \in \mathbb{R}[\mathbf{x}]_n,$$

which satisfies $s^2(0) = 2^n$ and $s^2(\mathbf{x}) = 0$ for all $0 \neq \mathbf{x} \in \mathbb{B}^n$. Clearly, we have:

$$\int_{\mathbb{B}^n} f(\mathbf{x}) s^2(\mathbf{x}) d\mu(\mathbf{x}) = f(0) = f_{\min} \text{ and } \int_{\mathbb{B}^n} s^2(\mathbf{x}) d\mu(\mathbf{x}) = 1,$$

and so $\mathrm{ub}(f)_n = f_{\min}$. □

The next lemma shows that this result is tight, by giving an example of polynomial f for which the bound $\mathrm{ub}(f)_r$ is exact only at order $r = n$.

LEMMA 7.19. *Let $f(\mathbf{x}) = |\mathbf{x}| = \mathbf{x}_1 + \ldots + \mathbf{x}_n$. Then $\mathrm{ub}(f)_r - f_{\min} > 0$ for all $r < n$.*

PROOF. Suppose not. That is, $\mathrm{ub}(f)_r = f_{\min} = 0$ for some $r \leq n-1$. As $f(\mathbf{x}) > 0 = f_{\min}$ for all $0 \neq \mathbf{x} \in \mathbb{B}^n$, this implies that there exists a polynomial $s \in \mathbb{R}[\mathbf{x}]_r$ such that s^2 is interpolating at 0, i.e. such that $s^2(0) = 1$ and $s^2(\mathbf{x}) = 0$ for all $0 \neq \mathbf{x} \in \mathbb{B}^n$. But then s is itself interpolating at 0 and has degree $r < n$, a contradiction. □

7.4. The harmonic constant

In this section we give a proof of Lemma 7.11, where we bound the sup-norm $\|p_k\|_\infty$ of the harmonic components p_k of a polynomial p by $\gamma_d \|p\|_\infty$ for some constant γ_d depending only on the degree d of p. The following definitions will be convenient.

Definition 7.20. *For $n \geq d \geq k \geq 0$ integers, we write:*

$$\gamma(\mathbb{B}^n)_{d,k} := \sup\{\|p_k\|_\infty : p = p_0 + p_1 + \cdots + p_d \in \mathcal{P}(\mathbb{B}^n)_d, \|p\|_\infty \leq 1\}, \text{ and}$$
$$\gamma(\mathbb{B}^n)_d := \max_{0\leq k\leq d} \gamma(\mathbb{B}^n)_{d,k}.$$

We are thus interested in finding a bound γ_d depending only on d such that:

$$\gamma_d \geq \gamma(\mathbb{B}^n)_d \text{ for all } n \in \mathbb{N}.$$

We will now show that in the computation of the parameter $\gamma(\mathbb{B}^n)_{d,k}$ we may restrict to feasible solutions p having strong structural properties. First, we show that we may assume that the sup-norm of the harmonic component p_k of p is attained at $0 \in \mathbb{B}^n$.

Lemma 7.21. *We have:*

$$\gamma(\mathbb{B}^n)_{d,k} = \sup_{p\in\mathcal{P}(\mathbb{B}^n)_d} \{p_k(0) : \|p\|_\infty \leq 1\} \tag{7.25}$$

Proof. Let p be a feasible solution for $\gamma(\mathbb{B}^n)_{d,k}$ and let $\mathbf{x} \in \mathbb{B}^n$ for which $p_k(\mathbf{x}) = \|p_k\|_\infty$ (after possibly replacing p by $-p$). Now choose $\sigma \in \mathrm{Aut}(\mathbb{B}^n)$ such that $\sigma(0) = \mathbf{x}$ and set $\widehat{p} = p \circ \sigma$. Clearly, $\widehat{p}$ is again a feasible solution for $\gamma(\mathbb{B}^n)_{d,k}$. Moreover, as H_k is invariant under $\mathrm{Aut}(\mathbb{B}^n)$, we have:

$$\|\widehat{p}_k\|_\infty = \widehat{p}_k(0) = (p \circ \sigma)_k(0) = (p_k \circ \sigma)(0) = \|p_k\|_\infty,$$

which shows the lemma. □

Next we show that we may in addition restrict to polynomials that are highly symmetric.

Lemma 7.22. *In the program* (7.25) *we may restrict the optimization to polynomials of the form:*

$$p(\mathbf{x}) = \sum_{i=0}^{d} \lambda_i \sum_{|a|=i} \chi_a(\mathbf{x}) = \sum_{i=0}^{d} \lambda_i \mathcal{K}_i^{(n)}(|\mathbf{x}|) \quad \text{where } \lambda_i \in \mathbb{R}.$$

Proof. Let p be a feasible solution to (7.25). Consider the following polynomial $\widehat{p}$ obtained as symmetrization of p under action of $\mathrm{St}(0)$, the set of automorphism of $\mathbb{B}^n$ corresponding to the coordinate permutations:

$$\widehat{p}(\mathbf{x}) = \frac{1}{|\mathrm{St}(0)|} \sum_{\sigma\in\mathrm{St}(0)} (p \circ \sigma)(\mathbf{x}).$$

Then $\|\widehat{p}\|_\infty \leq 1$ and $\widehat{p}_k(0) = p_k(0)$, so $\widehat{p}$ is still feasible for (7.25) and has the same objective value as p. Furthermore, for each i, $\widehat{p}_i$ is invariant under $\mathrm{St}(0)$, which implies that $\widehat{p}_i(\mathbf{x}) = \lambda_i X_i(\mathbf{x}) = \lambda_i \sum_{|a|=i} \chi_a(\mathbf{x}) = \lambda_i \mathcal{K}_i^{(n)}(|\mathbf{x}|)$ for some $\lambda_i \in \mathbb{R}$ (see (7.10)). □

A simple rescaling $\lambda_i \leftarrow \lambda_i \cdot \binom{n}{i}$ allows us to switch from $\mathcal{K}_i^{(n)}$ to $\widehat{\mathcal{K}}_i^n = \mathcal{K}_i^{(n)}/\binom{n}{i}$ and to obtain the following reformulation of $\gamma(\mathbb{B}^n)_{d,k}$ as a linear program.

LEMMA 7.23. *For any $n \geq d \geq k$ we have:*

$$\begin{aligned} \gamma(\mathbb{B}^n)_{d,k} = \max \quad & \lambda_k \\ s.t. \quad & -1 \leq \sum_{i=0}^{d} \lambda_i \widehat{\mathcal{K}}_i^n(t) \leq 1 \quad (t = 0, 1, \ldots, n). \end{aligned} \tag{7.26}$$

7.4.1. Limit functions. The idea now is to prove a bound on $\gamma(\mathbb{B}^n)_{d,k}$ which holds for fixed d and is independent of n. We will do this by considering 'the limit' of problem (7.26) as $n \to \infty$. For each $k \in \mathbb{N}$, we define the limit function:

$$\widehat{\mathcal{K}}_k^\infty(t) := \lim_{n\to\infty} \widehat{\mathcal{K}}_k^n(nt),$$

which, as shown in Lemma 7.25 below, is in fact a polynomial. We first present the polynomial $\widehat{\mathcal{K}}_k^\infty(t)$ for small k as an illustration.

EXAMPLE 7.24. *We have:*

$$\begin{aligned} \widehat{\mathcal{K}}_0^n(nt) = 1 &\implies \widehat{\mathcal{K}}_0^\infty(t) = 1, \\ \widehat{\mathcal{K}}_1^n(nt) = -2t + 1 &\implies \widehat{\mathcal{K}}_1^\infty(t) = -2t + 1, \\ \widehat{\mathcal{K}}_2^n(nt) = \frac{2n^2t^2 - 2n^2t + \binom{n}{2}}{\binom{n}{2}} &\implies \widehat{\mathcal{K}}_2^\infty(t) = 4t^2 - 4t + 1 = (1-2t)^2. \end{aligned}$$

LEMMA 7.25. *We have: $\widehat{\mathcal{K}}_k^\infty(t) = (1-2t)^k$ for all $k \in \mathbb{N}$.*

PROOF. The Krawtchouk polynomials satisfy the following three-term recurrence relation (see, e.g., [**MS83**]):

$$(k+1)\mathcal{K}_{k+1}^{(n)}(t) = (n-2t)\mathcal{K}_k^{(n)}(t) - (n-k+1)\mathcal{K}_{k-1}^{(n)}(t)$$

for $1 \leq k \leq n-1$. By evaluating the polynomials at nt we obtain:

$$\begin{aligned} & (k+1)\mathcal{K}_{k+1}^{(n)}(nt) = (n-2nt)\mathcal{K}_k^{(n)}(nt) - (n-k+1)\mathcal{K}_{k-1}^{(n)}(nt), \\ \implies \quad & (k+1)\binom{n}{k+1}\widehat{\mathcal{K}}_{k+1}^n(nt) = (n-2nt)\binom{n}{k}\widehat{\mathcal{K}}_k^n(nt) \\ & \qquad\qquad - (n-k+1)\binom{n}{k-1}\widehat{\mathcal{K}}_{k-1}^n(nt), \\ \implies \quad & \widehat{\mathcal{K}}_{k+1}^n(nt) = \frac{n(1-2t)}{(n-k)} \cdot \widehat{\mathcal{K}}_k^n(nt) - \frac{k}{n-k} \cdot \widehat{\mathcal{K}}_{k-1}^n(nt), \\ \implies \quad & \widehat{\mathcal{K}}_{k+1}^\infty(t) = (1-2t)\widehat{\mathcal{K}}_k^\infty(t). \end{aligned}$$

As $\widehat{\mathcal{K}}_0^\infty(t) = 1$ and $\widehat{\mathcal{K}}_1^\infty(t) = 1-2t$ we can conclude that indeed $\widehat{\mathcal{K}}_k^\infty(t) = (1-2t)^k$ for all $k \in \mathbb{N}$. □

Next, we show that solutions to (7.26) remain feasible after increasing the dimension n.

LEMMA 7.26. *Let $\lambda = (\lambda_0, \lambda_1, \ldots, \lambda_d)$ be a feasible solution to* (7.26) *for a certain $n \in \mathbb{N}$. Then it is also feasible to* (7.26) *for $n+1$ (and thus for any $n' \geq n+1$). Therefore, $\gamma(\mathbb{B}^{n+1})_{d,k} \geq \gamma(\mathbb{B}^n)_{d,k}$ for all $n \geq d \geq k$ and thus $\gamma(\mathbb{B}^{n+1})_d \geq \gamma(\mathbb{B}^n)_d$ for all $n \geq d$.*

PROOF. We may view $\mathbb{B}^n$ as a subset of $\mathbb{B}^{n+1}$ via the map $a \mapsto (a, 0)$, and analogously $\mathcal{P}(\mathbb{B}^n)$ as a subspace of $\mathcal{P}(\mathbb{B}^{n+1})$ via $\chi_a \mapsto \chi_{(a,0)}$. Now for $m, i \in \mathbb{N}$ we consider again the zonal spherical harmonic (7.10):

$$X_i^m = \sum_{|a|=i, a \in \mathbb{B}^m} \chi_a.$$

Consider the set $\mathrm{St}(0) \subseteq \mathrm{Aut}(\mathbb{B}^{n+1})$ of automorphisms fixing $0 \in \mathbb{B}^{n+1}$, i.e., the coordinate permutations arising from $\sigma \in \mathrm{Sym}(n+1)$. We will use the following identity:

$$\frac{1}{|\mathrm{St}(0)|} \sum_{\sigma \in \mathrm{St}(0)} \frac{X_i^n}{\binom{n}{i}} \circ \sigma = \frac{X_i^{n+1}}{\binom{n+1}{i}}. \tag{7.27}$$

To see that (7.27) holds note that its left hand side is equal to

$$\frac{1}{(n+1)!\binom{n}{i}} \sum_{\sigma \in \mathrm{Sym}(n+1)} \sum_{a \in \mathbb{B}^n, |a|=i} \chi_{(a,0)} \circ \sigma = \frac{1}{(n+1)!\binom{n}{i}} \sum_{b \in \mathbb{B}^{n+1}, |b|=i} N_b \chi_b,$$

where N_b denotes the number of pairs (σ, a) with $\sigma \in \mathrm{Sym}(n+1)$, $a \in \mathbb{B}^n$, $|a| = i$ such that $b = \sigma(a, 0)$. As there are $\binom{n}{i}$ choices for a and $i!(n+1-i)!$ choices for σ we have $N_b = \binom{n}{i} i!(n+1-i)!$ and thus (7.27) holds.

Assume λ is a feasible solution of (7.26) for a given value of n. Then, in view of (7.13), this means

$$\Big|\sum_{i=0}^d \lambda_i \cdot \frac{X_i^n(\mathbf{x})}{\binom{n}{i}}\Big| \leq 1 \quad \text{for all } x \in \mathbb{B}^n, \quad \text{and thus for all } \mathbf{x} \in \mathbb{B}^{n+1}.$$

Using (7.27) we obtain:

$$\begin{aligned}\Big|\sum_{i=0}^d \lambda_i \frac{X_i^{n+1}(\mathbf{x})}{\binom{n+1}{i}}\Big| &= \Big|\sum_{i=0}^d \lambda_i \cdot \frac{1}{|\mathrm{St}(0)|} \sum_{\sigma \in \mathrm{St}(0)} \frac{X_i^n(\sigma(\mathbf{x}))}{\binom{n}{i}}\Big| \\ &= \Big|\Big(\frac{1}{|\mathrm{St}(0)|} \sum_{\sigma \in \mathrm{St}(0)} \big(\sum_{i=0}^d \lambda_i \frac{X_i^n}{\binom{n}{i}}\big) \circ \sigma\Big)(\mathbf{x})\Big| \leq 1\end{aligned}$$

for all $\mathbf{x} \in \mathbb{B}^{n+1}$. Using (7.13) again, this shows that λ is a feasible solution of program (7.26) for $n+1$. □

Example 7.27. *To illustrate the identity* (7.27)*, we give a small example with* $n = i = 2$*. Consider:*

$$X_2^2 = \sum_{|a|=2, a\in\mathbb{B}^2} \chi_a = \chi_{11}.$$

The automorphisms in $\mathrm{St}(0) \subseteq \mathrm{Aut}(\mathbb{B}^3)$ *fixing* $0 \in \mathbb{B}^3$ *are the permutations of* x_1, x_2, x_3*. So we get:*

$$\begin{aligned}\frac{1}{|\mathrm{St}(0)|}\sum_{\sigma\in\mathrm{St}(0)} X_2^2 \circ \sigma &= \frac{1}{6}(\chi_{110} + \chi_{101} + \chi_{110} + \chi_{011} + \chi_{101} + \chi_{011}) \\ &= \frac{2}{6}(\chi_{110} + \chi_{101} + \chi_{011}) = \frac{1}{3}X_2^3,\end{aligned}$$

and indeed $\binom{2}{2}/\binom{3}{2} = 1/3$*.*

Lemma 7.28. *For* $d \geq k \in \mathbb{N}$*, define the program:*

$$\begin{aligned}\gamma_{d,k} := \max \quad & \lambda_k \\ s.t. \quad & -1 \leq \sum_{i=0}^{d} \lambda_i \widehat{\mathcal{K}}_i^{\infty}(x) \leq 1 \quad (x \in [0,1]).\end{aligned} \tag{7.28}$$

Then, for any $n \geq d$*, we have:* $\gamma(\mathbb{B}^n)_{d,k} \leq \gamma_{d,k}$*.*

Proof. Let λ be a feasible solution to (7.26) for (n, d, k). We show that λ is feasible for (7.28). For this fix $t \in [0,1] \cap \mathbb{Q}$. Then there exists a sequence of integers $(n_j)_j \to \infty$ such that $n_j \geq n$ and $tn_j \in [0, n_j]$ is integer for each $j \in \mathbb{N}$. As $n_j \geq n$, we know from Lemma 7.26 that λ is also a feasible solution of program (7.26) for (n_j, d, k). Hence, since $n_j t \in [0 : n_j]$ we obtain

$$|\sum_{i=0}^{d} \lambda_i \widehat{\mathcal{K}}_i^{n_j}(n_j t)| \leq 1 \quad \forall j \in \mathbb{N}.$$

But this immediately gives:

$$|\sum_{i=0}^{d} \lambda_i \widehat{\mathcal{K}}_i^{\infty}(t)| = \lim_{j\to\infty} |\sum_{i=0}^{d} \lambda_i \widehat{\mathcal{K}}_i^{n_j}(n_j t)| \leq 1. \tag{7.29}$$

As $[0,1] \cap \mathbb{Q}$ lies dense in $[0,1]$ (and the $\widehat{\mathcal{K}}_i^{\infty}$'s are continuous) we may conclude that (7.29) holds for all $t \in [0,1]$. This shows that λ is feasible for (7.28) and we thus have $\gamma(\mathbb{B}^n)_{d,k} \leq \gamma_{d,k}$, as desired. □

It remains now to compute the optimum solution to the program (7.28). In light of Lemma 7.25, and after a change of variables, this program may be reformulated as:

$$\begin{aligned}\max \quad & |\lambda_k| \\ s.t. \quad & -1 \leq \sum_{i=0}^{d} \lambda_i x^i \leq 1 \quad (x \in [-1,1]).\end{aligned} \tag{7.30}$$

In other words, we are tasked with finding a polynomial $p(x)$ of degree d satisfying $|p(\mathbf{x})| \leq 1$ for all $x \in [-1, 1]$, whose k-th coefficient is as large as possible in absolute value. This is a classical extremal problem solved by V. Markov.

THEOREM 7.29 (see, e.g., Theorem 7, pp. 53 in [**Nat64**]). *For $m \in \mathbb{N}$, let $T_m(x) = \sum_{i=0}^{m} t_{m,i} x^i$ be the Chebyshev polynomial of degree m. Then the optimum solution λ to* (7.30) *is given by:*

$$\sum_{i=0}^{d} \lambda_i x^i = \begin{cases} T_d(x) & \text{if } k \equiv d \mod 2, \\ T_{d-1}(x) & \text{if } k \not\equiv d \mod 2. \end{cases}$$

In particular, $\gamma_{d,k}$ is equal to $|t_{d,k}|$ (resp. $|t_{d-1,k}|$).

As the coefficients of the Chebyshev polynomials are known explicitely, Theorem 7.29 allows us to give exact values of the constant γ_d appearing in our main results (see Table 7.1). Using the following identity:

$$\sum_{i=0}^{d} |t_{d,i}| = \frac{1}{2}(1+\sqrt{2})^d + \frac{1}{2}(1-\sqrt{2})^d \leq (1+\sqrt{2})^d,$$

we are also able to concretely estimate:

$$\gamma_d \leq \max_{k \leq d} \gamma_{d,k} \leq (1+\sqrt{2})^d.$$

7.5. The q-ary cube

In this section, we indicate how our results for the binary cube $\mathbb{B}^n$ may be extended to the q-ary cube $(\mathbb{Z}/q\mathbb{Z})^n = \{0, 1, \ldots, q-1\}^n$ when $q > 2$ is a fixed integer. Here $\mathbb{Z}/q\mathbb{Z}$ denotes the cyclic group of order q, so that $(\mathbb{Z}/q\mathbb{Z})^n = \mathbb{B}^n$ when $q = 2$. The lower bound $\text{lb}(f)_r$ for the minimum of a polynomial f over $(\mathbb{Z}/q\mathbb{Z})^n$ is defined analogously to the case $q = 2$; namely we set

$$\text{lb}(f)_r := \sup_{\lambda \in \mathbb{R}} \{\lambda : f(\mathbf{x}) - \lambda \text{ is sos of degree at most } 2r \text{ on } (\mathbb{Z}/q\mathbb{Z})^n\},$$

where the condition means that $f(\mathbf{x}) - \lambda$ agrees with a sum of squares $s \in \Sigma[\mathbf{x}]_{2r}$ for all $x \in (\mathbb{Z}/q\mathbb{Z})^n$ or, alternatively, that $f - \lambda - s$ belongs to the ideal generated by the polynomials $\mathbf{x}_i(\mathbf{x}_i - 1) \ldots (\mathbf{x}_i - q + 1)$ for $i \in [n]$. Similarly, the upper bound $\text{ub}(f)_r$ is defined as in (7.3) after equipping $(\mathbb{Z}/q\mathbb{Z})^n$ with the uniform measure μ. The parameters $\text{lb}(f)_r$ and $\text{ub}(f)_r$ may again be computed by solving a semidefinite program of size polynomial in n for fixed $r, q \in \mathbb{N}$, see [**Lau07a**].

As before $d(\mathbf{x}, \mathbf{y})$ denotes the Hamming distance and $|\mathbf{x}|$ denotes the Hamming weight (number of nonzero components). Note that, for $\mathbf{x}, \mathbf{y} \in (\mathbb{Z}/q\mathbb{Z})^n$, $d(\mathbf{x}, \mathbf{y})$ can again be expressed as a polynomial in $\mathbf{x}, \mathbf{y}$, with degree $q - 1$ in each of $\mathbf{x}$ and $\mathbf{y}$.

We will prove Theorem 7.32 below, which can be seen as an analog of Corollary 7.4 for $(\mathbb{Z}/q\mathbb{Z})^n$. The general structure of the proof is identical to that of the case $q = 2$. We therefore only give generalizations of arguments as necessary. For reasons that will become clear later, it is most convenient to consider the sum-of-squares bound $\mathrm{lb}(f)_r$ on the minimum $f_{\min}$ of a polynomial f with degree at most $(q-1)d$ over $(\mathbb{Z}/q\mathbb{Z})^n$, where $d \le n$ is fixed.

7.5.1. Fourier analysis on $(\mathbb{Z}/q\mathbb{Z})^n$ and Krawtchouk polynomials. Consider the space

$$\mathcal{P}((\mathbb{Z}/q\mathbb{Z})^n) := \mathbb{C}[\mathbf{x}]/(\mathbf{x}_i(\mathbf{x}_i - 1)\dots(\mathbf{x}_i - q + 1) : i \in [n])$$

consisting of the polynomials on $(\mathbb{Z}/q\mathbb{Z})^n$ with complex coefficients. We equip $\mathcal{P}((\mathbb{Z}/q\mathbb{Z})^n)$ with its natural *complex* inner product:

$$\langle f, g\rangle_\mu = \int_{(\mathbb{Z}/q\mathbb{Z})^n} f(\mathbf{x})\overline{g(\mathbf{x})}d\mu(\mathbf{x}) = \frac{1}{q^n}\sum_{\mathbf{x}\in(\mathbb{Z}/q\mathbb{Z})^n} f(\mathbf{x})\overline{g(\mathbf{x})},$$

where μ is the uniform measure on $(\mathbb{Z}/q\mathbb{Z})^n$. The space $\mathcal{P}((\mathbb{Z}/q\mathbb{Z})^n)$ has dimension $|(\mathbb{Z}/q\mathbb{Z})^n| = q^n$ over $\mathbb{C}$ and it is spanned by the polynomials of degree up to $(q-1)n$. The reason we now need to work with polynomials with complex coefficients is that the characters have complex coefficients when $q > 2$.

Let $\psi = e^{2\pi i/q}$ be a primitive q-th root of unity. For $a \in (\mathbb{Z}/q\mathbb{Z})^n$, the associated *character* $\chi_a \in \mathcal{P}((\mathbb{Z}/q\mathbb{Z})^n)$ is defined by:

$$\chi_a(\mathbf{x}) = \psi^{a\cdot\mathbf{x}} \quad (\mathbf{x} \in (\mathbb{Z}/q\mathbb{Z})^n).$$

So (7.7) is indeed the special case of this definition when $q = 2$. The set of all characters $\{\chi_a : a \in (\mathbb{Z}/q\mathbb{Z})^n\}$ forms an orthogonal basis for $\mathcal{P}((\mathbb{Z}/q\mathbb{Z})^n)$ w.r.t. the above inner product $\langle\cdot,\cdot\rangle_\mu$. A character χ_a can be written as a polynomial of degree $(q-1)\cdot|a|$ on $(\mathbb{Z}/q\mathbb{Z})^n$, i.e., we have $\chi_a \in \mathcal{P}((\mathbb{Z}/q\mathbb{Z})^n)_{(q-1)|a|}$ for all $a \in (\mathbb{Z}/q\mathbb{Z})^n$.

As before, we have the direct sum decomposition into pairwise orthogonal subspaces:

$$\mathcal{P}((\mathbb{Z}/q\mathbb{Z})^n) = H_0 \perp H_1 \perp \dots \perp H_n,$$

where H_i is spanned by the set $\{\chi_a : |a| = i\}$ and $H_i \subseteq \mathbb{R}[\mathbf{x}]_{(q-1)i}$. The components H_i are invariant and irreducible under the action of $\mathrm{Aut}((\mathbb{Z}/q\mathbb{Z})^n)$, which is generated by the coordinate permutations and the action of $\mathrm{Sym}(q)$ on individual coordinates. Hence any $p \in \mathcal{P}((\mathbb{Z}/q\mathbb{Z})^n)$ of degree at most $(q-1)d$ can be (uniquely) decomposed as:

$$p = p_0 + p_1 + \dots + p_d \quad (p_i \in H_i).$$

As before $\mathrm{St}(0) \subseteq \mathrm{Aut}((\mathbb{Z}/q\mathbb{Z})^n)$ denotes the stabilizer of $0 \in (\mathbb{Z}/q\mathbb{Z})^n$, which is generated by the coordinate permutations and the permutations in $\mathrm{Sym}(q)$ fixing 0 in $\{0, 1, \dots, q-1\}$ at any individual coordinate. We note for later reference that the subspace of H_i invariant under action of $\mathrm{St}(0)$ is of dimension one, and is spanned by the zonal spherical function:

$$X_i = \sum_{|a|=i} \chi_a \in H_i. \tag{7.31}$$

The Krawtchouk polynomials introduced in Section 7.1 have the following generalization in the q-ary setting:

$$\mathcal{K}_k^{(n)}(x) = \mathcal{K}_{k,q}^{(n)}(x) := \sum_{i=0}^{k} (-1)^i (q-1)^{k-i} \binom{x}{i} \binom{n-x}{k-i}.$$

Analogously to relation (7.12), the Krawtchouk polynomials $\mathcal{K}_k^{(n)}$ $(0 \leq k \leq n)$ are pairwise orthogonal w.r.t. the discrete measure ω on $[0:n]$ given by:

$$\omega(x) = \frac{1}{q^n} \sum_{x=0}^{n} w(x)\delta_x, \text{ with } w(x) := (q-1)^x \binom{n}{x}.$$

To be precise, we have:

$$\sum_{x=0}^{n} \mathcal{K}_k^{(n)}(x)\mathcal{K}_{k'}^{(n)}(x)(q-1)^x \binom{n}{x} = \delta_{k,k'}(q-1)^k \binom{n}{k}.$$

As $\mathcal{K}_k^{(n)}(0) = (q-1)^k \binom{n}{k} = \|\mathcal{K}_k^{(n)}\|_\omega^2$, we may normalize $\mathcal{K}_k^{(n)}$ by setting:

$$\widehat{\mathcal{K}}_k^n(x) := \mathcal{K}_k^{(n)}(x)/\mathcal{K}_k^{(n)}(0) = \mathcal{K}_k^{(n)}(x)/\|\mathcal{K}_k^{(n)}\|_\omega^2,$$

so that $\widehat{\mathcal{K}}_k^n$ satisfies $\max_{x=0}^n \widehat{\mathcal{K}}_k^n(x) = \widehat{\mathcal{K}}_k^n(0) = 1$ (cf. (7.15)).

We have the following connection (cf. (7.14)) between the characters and the Krawtchouk polynomials:

$$\sum_{a \in (\mathbb{Z}/q\mathbb{Z})^n : |a|=k} \chi_a(\mathbf{x}) = \mathcal{K}_k^{(n)}(i) \quad \text{for } x \in (\mathbb{Z}/q\mathbb{Z})^n \text{ with } |\mathbf{x}| = i. \tag{7.32}$$

Note that for all $a, \mathbf{x}, \mathbf{y} \in (\mathbb{Z}/q\mathbb{Z})^n$, we have:

$$\chi_a^{-1}(\mathbf{x}) = \overline{\chi_a(\mathbf{x})} = \chi_a(-\mathbf{x}), \quad \chi_a(\mathbf{x})\chi_a(\mathbf{y}) = \chi_a(\mathbf{x}+\mathbf{y}).$$

Hence, for any $\mathbf{x}, \mathbf{y} \in (\mathbb{Z}/q\mathbb{Z})^n$, we also have (cf. (7.13)):

$$\sum_{a \in (\mathbb{Z}/q\mathbb{Z})^n : |a|=k} \chi_a(\mathbf{x})\overline{\chi_a(\mathbf{y})} = \sum_{a \in (\mathbb{Z}/q\mathbb{Z})^n : |a|=k} \chi_a(\mathbf{x}-\mathbf{y}) = \mathcal{K}_k^{(n)}(i)$$

when $d(\mathbf{x}, \mathbf{y}) = |\mathbf{x} - \mathbf{y}| = i$.

7.5.2. Invariant kernels. In analogy to the binary case $q = 2$, for a degree r univariate polynomial $u \in \mathbb{R}[x]_r$ we define the associated polynomial kernel $\mathrm{K}(\mathbf{x}, \mathbf{y}) := u^2(d(\mathbf{x}, \mathbf{y}))$ $(x, \mathbf{y} \in (\mathbb{Z}/q\mathbb{Z})^n)$ and the associated kernel operator:

$$\mathbf{K}p(\mathbf{x}) = \int_{(\mathbb{Z}/q\mathbb{Z})^n} \overline{p(\mathbf{y})}\mathrm{K}(\mathbf{x}, \mathbf{y})d\mu(\mathbf{y}) = \frac{1}{q^n} \sum_{\mathbf{y} \in (\mathbb{Z}/q\mathbb{Z})^n} \overline{p(\mathbf{y})}\mathrm{K}(\mathbf{x}, \mathbf{y}) \quad (p \in \mathcal{P}((\mathbb{Z}/q\mathbb{Z})^n)).$$

Note that $\mathrm{K}(\mathbf{x}, \mathbf{y})$ is a polynomial on $(\mathbb{Z}/q\mathbb{Z})^n$ with degree $2r(q-1)$ in each of $\mathbf{x}$ and $\mathbf{y}$. Let us decompose the univariate polynomial $u(x)^2$ in the Krawtchouk basis as

$$u(x)^2 = \sum_{i=0}^{2r} \lambda_i \mathcal{K}_i^{(n)}(x).$$

Then the kernel operator $\mathbf{K}$ acts as follows on characters: for $z \in (\mathbb{Z}/q\mathbb{Z})^n$,

$$\mathbf{K}\chi_z = \lambda_{|z|}\chi_z,$$

which can be seen by retracing the proof of Theorem 7.9, and we obtain the Funk-Hecke formula (recall (7.19)): for any polynomial $p \in \mathcal{P}((\mathbb{Z}/q\mathbb{Z})^n)_{(q-1)d}$ with Harmonic decomposition $p = p_0 + \ldots + p_d$,

$$\mathbf{K}p = \lambda_0 p_0 + \cdots + \lambda_d p_d.$$

7.5.3. Performing the analysis. It remains to find a univariate polynomial $u \in \mathbb{R}[x]$ of degree r with $u^2(x) = \sum_{i=0}^{2r} \lambda_i \mathcal{K}_i^{(n)}(x)$ for which $\lambda_0 = 1$ and the other scalars λ_i are close to 1. As before (cf. (7.20)), we have:

$$\lambda_i = \langle \mathcal{K}_i^{(n)}, u^2 \rangle_\omega / \|\mathcal{K}_i^{(n)}\|_\omega^2 = \langle \widehat{\mathcal{K}}_i^n, u^2 \rangle_\omega.$$

So we would like to minimize $\sum_{i=1}^{2r}(1-\lambda_i)$. We are therefore interested in the inner Lasserre hierarchy applied to the minimization of the function $g(x) = d - \sum_{i=0}^{d} \widehat{\mathcal{K}}_i^n(x)$ on the set $[0:n]$ (equipped with the measure ω from (7.5.1)). We show first that this function g again has a nice linear upper estimator.

LEMMA 7.30. *We have:*

$$\begin{aligned} |\widehat{\mathcal{K}}_k^n(x) - \widehat{\mathcal{K}}_k^n(x+1)| &\le \frac{2k}{n}, \qquad && (x = 0, 1, \ldots, n-1) \\ |\widehat{\mathcal{K}}_k^n(x) - 1| &\le \frac{2k}{n} \cdot x \qquad && (x = 0, 1, \ldots, n) \end{aligned} \tag{7.33}$$

for all $k \le n$.

PROOF. The proof is almost identical to that of Lemma 7.8. Let $x \in [0:n-1]$ and $0 < k \le d$. Consider the elements $1^x 0^{n-x}, 1^{x+1}0^{n-x-1} \in (\mathbb{Z}/q\mathbb{Z})^n$ from Lemma 7.6. Then we have:

$$\begin{aligned} |\mathcal{K}_k^{(n)}(x) - \mathcal{K}_k^{(n)}(x+1)| &\overset{(7.32)}{=} |\sum_{|a|=k} \chi_a(1^x 0^{n-x}) - \chi_a(1^{x+1}0^{n-x-1})| \\ &\le 2 \cdot \#\{a \in (\mathbb{Z}/q\mathbb{Z})^n : |a| = k, a_{x+1} \neq 0\} \\ &= 2 \cdot (q-1)^k \cdot \binom{n-1}{k-1}, \end{aligned}$$

where for the inequality we note that $\chi_a(1^x 0^{n-x}) = \chi_a(1^{x+1}0^{n-x-1})$ if $a_{x+1} = 0$. As $\mathcal{K}_k^{(n)}(0) = (q-1)^k \binom{n}{k}$, this implies that:

$$|\widehat{\mathcal{K}}_k^n(x) - \widehat{\mathcal{K}}_k^n(x+1)| \le 2 \cdot \binom{n-1}{k-1} / \binom{n}{k} = \frac{2k}{n}.$$

This shows the first inequality of (7.33). The second inequality follows using the triangle inequality, a telescope summation argument and the fact that $\widehat{\mathcal{K}}_k^n(0) = 1$. □

From Lemma 7.30 we obtain that the function $g(x) = d - \sum_{i=0}^d \widehat{\mathcal{K}}_i^n(x)$ admits the following linear upper estimator: $g(x) \leq d(d+1) \cdot (x/n)$ for $x \in [0:n]$. Now the same arguments as used for the case $q = 2$ enable us to conclude:

$$\mathrm{ub}(f)_{(q-1)r} - f_{\min} \leq C_d \cdot \xi_{r+1,q}^n / n$$

and, when $d(d+1)\xi_{r+1,q}^n/n \leq 1/2$,

$$f_{\min} - \mathrm{lb}(f)_{(q-1)r} \leq 2C_d \cdot \xi_{r+1,q}^n/n.$$

Here C_d is a constant depending only on d and $\xi_{r+1,q}^n$ is the least root of the Krawtchouk polynomial $\mathcal{K}_{r+1,q}^{(n)}$. Note that as the kernel $\mathrm{K}(\mathbf{x}, \mathbf{y}) = u^2(d(\mathbf{x}, \mathbf{y}))$ is of degree $2(q-1)r$ in $\mathbf{x}$ (and $\mathbf{y}$), we are only able to analyze the corresponding levels $(q-1)r$ of the hierarchies. We come back below to the question on how to show the existence of the above constant C_d.

But first we finish the analysis. Having shown analogs of Theorem 7.1 and Theorem 7.3 in this setting, it remains to state the following more general version of Theorem 7.2, giving information about the smallest roots of the q-ary Krawtchouk polynomials.

THEOREM 7.31 ([**Lev98**], Section 5). *Fix $t \in [0, \frac{q-1}{q}]$. Then the smallest roots $\xi_{r,q}^n$ of the q-ary Krawtchouk polynomials $\mathcal{K}_{r,q}^{(n)}$ satisfy:*

$$\lim_{r/n \to t} \xi_{r,q}^n / n = \varphi_q(t) := \frac{q-1}{q} - \left(\frac{q-2}{q} \cdot t + \frac{2}{q}\sqrt{(q-1)t(1-t)} \right).$$

Here the above limit means that, for any sequences $(n_j)_j$ and $(r_j)_j$ of integers such that $\lim_{j\to\infty} n_j = \infty$ and $\lim_{j\to\infty} r_j/n_j = t$, we have $\lim_{j\to\infty} \xi_{r_j,q}^{n_j}/n_j = \varphi_q(t)$.

Note that for $q = 2$ we have $\varphi_q(t) = \frac{1}{2} - \sqrt{t(1-t)}$, which is the function $\varphi(t)$ from (7.5). To avoid technical details we only quote in Theorem 7.31 the asymptotic analog of Theorem 7.2 (and not the exact bound on the root $\xi_{r,q}^n$ for any n). Therefore we have shown the following q-analog of Corollary 7.4.

THEOREM 7.32. *Fix $d \leq n$ and for $n, r \in \mathbb{N}$ write*

$$E_{(r)}(n) := \sup_{f \in \mathbb{R}[\mathbf{x}]_{(q-1)d}} \big\{ f_{\min} - \mathrm{lb}(f)_r : \|f\|_\infty = 1 \big\},$$

$$E^{(r)}(n) := \sup_{f \in \mathbb{R}[\mathbf{x}]_{(q-1)d}} \big\{ \mathrm{ub}(f)_r - f_{\min} : \|f\|_\infty = 1 \big\}.$$

There exists a constant $C_d > 0$ (depending also on q) such that, for any $t \in [0, \frac{q-1}{q}]$, we have:

$$\lim_{r/n \to t} E^{((q-1)r)}(n) \leq C_d \cdot \varphi_q(t)$$

and, if $d(d+1) \cdot \varphi_q(t) \leq 1/2$, then we also have:

$$\lim_{r/n \to t} E_{((q-1)r)}(n) \leq 2 \cdot C_d \cdot \varphi_q(t).$$

Here $\varphi_q(t)$ is the function defined in (7.31). Recall that the limit notation $r/n \to t$ means that the claimed convergence holds for any sequences $(n_j)_j$ and $(r_j)_j$ of integers such that $\lim_{j\to\infty} n_j = \infty$ and $\lim_{j\to\infty} r_j/n_j = t$.

For reference, the function $\varphi_q(t)$ is shown for several values of q in Figure 7.1.

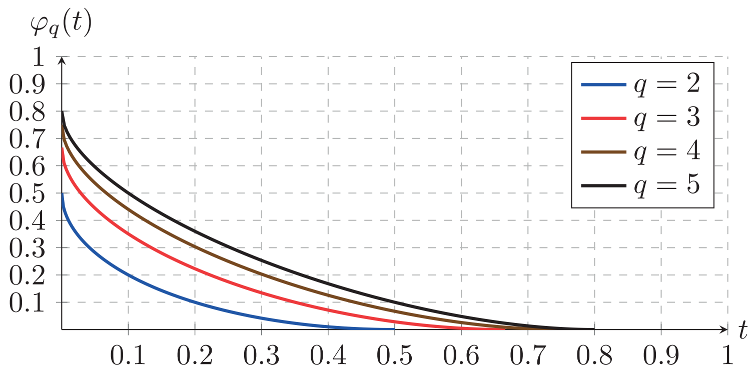

FIGURE 7.1. The function $\varphi_q(t)$ for several values of q. Note that the case $q = 2$ corresponds to the function $\varphi(t)$ of (7.5).

7.5.4. A generalization of Lemma 7.11. The arguments above omit a generalization of Lemma 7.11, which is instrumental to show the existence of the constant C_d claimed above. In other words, we still need to show that if $p : (\mathbb{Z}/q\mathbb{Z})^n \to \mathbb{R}$ is a polynomial of degree $(q-1)d$ on $(\mathbb{Z}/q\mathbb{Z})^n$ with harmonic decomposition $p = p_0 + \ldots + p_d$, there then exists a constant $\gamma_d > 0$ (independent of n) such that:

$$\|p_i\|_\infty \leq \gamma_d \|p\|_\infty \text{ for all } 0 \leq i \leq d.$$

Then, as in the binary case, we may set $C_d = d(d+1)\gamma_d$. The proof given in Section 7.4 for the case $q = 2$ applies almost directly to the general case, and we only generalize certain steps as required. So consider again the parameters:

$$\gamma((\mathbb{Z}/q\mathbb{Z})^n)_{d,k} := \sup\{\|p_k\|_\infty : p = p_0 + p_1 + \cdots + p_d \in \mathbb{R}[\mathbf{x}]_{(q-1)d}, \|p\|_\infty \leq 1\},$$
$$\gamma((\mathbb{Z}/q\mathbb{Z})^n)_d := \max_{0 \leq k \leq d} \gamma((\mathbb{Z}/q\mathbb{Z})^n)_{d,k}.$$

Lemmas 7.21 and 7.22, which show that the optimum solution p to $\gamma((\mathbb{Z}/q\mathbb{Z})^n)_{d,k}$ may be assumed to be invariant under $\mathrm{St}(0) \subseteq \mathrm{Aut}((\mathbb{Z}/q\mathbb{Z})^n)$, clearly apply to the case $q > 2$ as well. That is to say, we may assume p is of the form[2]:

$$p(\mathbf{x}) = \sum_{i=0}^{d} \lambda_i X_i(\mathbf{x}) \quad (\lambda_i \in \mathbb{R})$$

where $X_i = \sum_{|a|=i} \chi_a \in H_i$ is the zonal spherical function of degree $(q-1)i$ (cf. (7.31) and (7.10)). Using (7.32), we obtain a reformulation of $\gamma((\mathbb{Z}/q\mathbb{Z})^n)_{d,k}$ as an LP (cf. (7.26)):

$$\begin{aligned}\gamma((\mathbb{Z}/q\mathbb{Z})^n)_{d,k} = \max \quad & \lambda_k \\ s.t. \quad & -1 \leq \sum_{i=0}^{d} \lambda_i \widehat{\mathcal{K}}_{i,q}^n(x) \leq 1 \quad (x = 0, 1, \ldots, n).\end{aligned}$$

For $k \in \mathbb{N}$, let $\widehat{\mathcal{K}}_k^\infty(x) := \lim_{n\to\infty} \widehat{\mathcal{K}}_k^n(nx) = \left(1 - \frac{q}{q-1}x\right)^k$ and consider the program (cf. (7.28)):

$$\begin{aligned}\gamma_{d,k} := \max \quad & \lambda_k \\ s.t. \quad & -1 \leq \sum_{i=0}^{d} \lambda_i \widehat{\mathcal{K}}_i^\infty(x) \leq 1 \quad (x \in [0, 1]).\end{aligned}$$

As before, we have $\gamma((\mathbb{Z}/q\mathbb{Z})^n)_{d,k} \leq \gamma_{d,k}$, noting that (the proofs of) Lemma 7.26 and Lemma 7.28 may be applied directly to the case $q > 2$. From there, it suffices to show $\gamma_{d,k} < \infty$, which can be argued in an analogous way to the case $q = 2$.

7.6. Discussion

Using the polynomial kernel method introduced in Chapter 6, we have shown a theoretical guarantee on the quality of the sum-of-squares hierarchy $\mathrm{lb}(f)_r \leq f_{\min}$ for approximating the minimum of a polynomial f of degree d over the binary cube $\mathbb{B}^n$. As far as we are aware, this is the first such analysis that applies to values of r smaller than $(n+d)/2$, i.e., when the hierarchy is not exact. Additionally as we explained in Section 6.2, our guarantee may be extended to the measure-based hierarchy of bounds $\mathrm{ub}(f)_r \geq f_{\min}$. Our result may therefore also be interpreted as bounding the range $\mathrm{ub}(f)_r - \mathrm{lb}(f)_r$. Our analysis also applies to polynomial optimization over the cube $\{\pm 1\}^n$ (by a simple change of variables) and over the q-ary cube.

[2]Note that as p is assumed to be real-valued, the coefficients λ_i must be real. Indeed, for each $a \in (\mathbb{Z}/q\mathbb{Z})^n$, we have $\langle p, \chi_a \rangle_\mu = \lambda_{|a|} \|\chi_a\|^2 = \lambda_{|a|} \|\chi_a^{-1}\|^2 = \langle p, \overline{\chi_a} \rangle_\mu = \overline{\langle p, \chi_a \rangle_\mu}$.

Analysis for small values of r. A limitation of Theorem 7.1 is that the analysis of $\mathrm{lb}(f)_r$ applies only for choices of d, r, n satisfying $d(d+1)\xi^n_{r+1} \leq 1/2$. One may partially avoid this limitation by proving a slightly sharper version of Lemma 7.13, showing instead that $\Lambda \leq \tilde{\Lambda}/(1-\tilde{\Lambda})$, assuming now only that $\tilde{\Lambda} \leq 1$. Indeed, Lemma 7.13 is a special case of this result, assuming that $\tilde{\Lambda} \leq 1/2$ to obtain $\Lambda \leq 2\tilde{\Lambda}$. Nevertheless, our methods exclude values of r outside of the regime $r = \Omega(n)$.

The constant γ_d. The strength of our results depends in large part on the size of the constant C_d appearing in Theorem 7.1 and Theorem 7.3, where we may set $C_d = d(d+1)\gamma_d$. Recall that γ_d is defined in Lemma 7.11 as a constant for which $\|p_k\|_\infty \leq \gamma_d \|p\|_\infty$ for any polynomial $p = p_0 + p_1 + \ldots + p_d$ of degree d and $k \leq d$ on $\mathbb{B}^n$, independently of the dimension n. In Section 7.4 we have shown the existence of such a constant. Furthermore, we have shown there that we may choose $\gamma_d \leq (1+\sqrt{2})^d$, and have given an explicit expression for the smallest possible value of γ_d in terms of the coefficients of Chebyshev polynomials. Table 7.1 lists these values for small d.

d	1	2	3	4	5	6	7	8	9	10
γ_d	1	2	4	8	20	48	112	256	576	1280

TABLE 7.1. Values of the constant γ_d.

Computing extremal roots of Krawtchouk polynomials. Although Theorem 7.2 provides only an asymptotic bound on the least root ξ^n_r of $\mathcal{K}^{(n)}_r$, it should be noted that ξ^n_r can be computed explicitely for small values of r, n, thus allowing for a concrete estimate of the error of both Lasserre hierarchies via Theorem 7.1 and Theorem 7.3, respectively. Indeed, as we have seen, the root ξ^n_{r+1} is equal to the smallest eigenvalue of the $(r+1) \times (r+1)$ matrix J (aka Jacobi matrix), whose entries are given by $J_{i,j} = \langle x\widehat{\mathcal{K}}^n_i(x), \widehat{\mathcal{K}}^n_j(x)\rangle_\omega$ for $i, j \in \{0, 1, \ldots, r\}$. See also [**Sze75**] for details.

Matrix-valued polynomials. The results of this chapter carry over to the setting of *matrix-valued* polynomials. Indeed, the proof technique we use may be applied there rather straightforwardly. This was already noted in [**FF21**] in the context of polynomial optimization on the hypersphere S^{n-1}. As matrix-valued polynomials fall outside the scope of this thesis, we refer the reader to our paper [**SL21b**] (on which this chapter is based) for details.

Acknowledgments. We wish to thank Sven Polak and Pepijn Roos Hoefgeest for several useful discussions.

CHAPTER 8

Application: The unit ball and standard simplex

Als je het opschrijft, staat het meteen op papier ook.

Gerard Reve

This chapter is based on my work [**Slo22**].

In this chapter, we apply the method of Chapter 6 to obtain convergence rates for sum-of-squares hierarchies on the unit ball B^n and on the standard simplex Δ^n. Let us briefly recall the setup. Let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree d. We consider the polynomial optimization problem:

$$f_{\min} := \min_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x}),$$

where $\mathbf{X} = B^n$ or $\mathbf{X} = \Delta^n$. The unit ball and the standard simplex are both semialgebraic sets, with description:

$$B^n = \{\mathbf{x} \in \mathbb{R}^n : 1 - \|\mathbf{x}\|^2 = 1 - \sum_{i=1}^n \mathbf{x}_i^2 \geq 0\},$$

$$\Delta^n = \{\mathbf{x} \in \mathbb{R}^n : 1 - \sum_{i=1}^n \mathbf{x}_i \geq 0,\ \mathbf{x}_i \geq 0 \quad (1 \leq i \leq n)\}.$$

We may thus define lower bounds $\mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r, \mathrm{lb}(f, \mathcal{T}(\mathbf{X}))_r$ on the minimum $f_{\min}$ using the (truncated) quadratic module $\mathcal{Q}(\mathbf{X})_{2r}$ and preordering $\mathcal{T}(\mathbf{X})_{2r}$, respectively. See Section 2.1 for the definitions. We note that since B^n is defined using only a single inequality constraint, we have $\mathcal{Q}(B^n)_{2r} = \mathcal{T}(B^n)_{2r}$ for each $r \in \mathbb{N}$, and the Putinar- and Schmüdgen-type bounds for B^n thus coincide:

$$\mathrm{lb}(f, \mathcal{Q}(B^n))_r = \mathrm{lb}(f, \mathcal{T}(B^n))_r.$$

If $\mathbf{X}$ is equipped with a finite Borel measure μ, we also have corresponding hierarchies of *upper* bounds on $f_{\min}$ (see Section 2.1), which are given by:

$$\mathrm{ub}(f, \mathcal{Q}(\mathbf{X}), \mu)_r = \inf_{q \in \mathcal{Q}(\mathbf{X})_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mu(\mathbf{x}) : \int_{\mathbf{X}} q(\mathbf{x}) d\mu(\mathbf{x}) = 1 \right\},$$

$$\mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r = \inf_{q \in \mathcal{T}(\mathbf{X})_{2r}} \left\{ \int_{\mathbf{X}} f(\mathbf{x}) q(\mathbf{x}) d\mu(\mathbf{x}) : \int_{\mathbf{X}} q(\mathbf{x}) d\mu(\mathbf{x}) = 1 \right\} \quad (8.1)$$

Main results. The main contribution of this chapter is to show a convergence rate in $O(1/r^2)$ of the lower bounds $\text{lb}(f, \mathcal{T}(\mathbf{X}))_r$ to the global minimum $f_{\min}$ of a polynomial f on the unit ball $\mathbf{X} = B^n$ or on the standard simplex $\mathbf{X} = \Delta^n$. These rates match the best-known rates for the hypersphere S^{n-1} of [**FF21**], see also Table 2.1. For the unit ball, no (specialized) rates were known before. For the simplex, we improve upon the previously best known bound in $O(1/r)$ due to Kirschner & de Klerk [**KdK21**].

THEOREM 8.1. *Let* $\mathbf{X} = B^n = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|^2 \leq 1\}$ *be the n-dimensional unit ball and let* $f \in \mathbb{R}[\mathbf{x}]$ *be a polynomial of degree d. Then for any* $r \geq 2nd$, *the lower bound* $\text{lb}(f, \mathcal{T}(\mathbf{X}))_r$ *for the minimization of* f *over* B^n *satisfies:*

$$f_{\min} - \text{lb}(f, \mathcal{T}(\mathbf{X}))_r \leq \frac{C_B(n,d)}{r^2} \cdot (f_{\max} - f_{\min}).$$

Here, $C_B(n,d)$ *is a constant depending only on* n, d. *This constant depends polynomially on n (for fixed d) and polynomially on d (for fixed n). See relation* (8.28) *for details.*

THEOREM 8.2. *Let* $\mathbf{X} = \Delta^n = \{\mathbf{x} \in \mathbb{R}^n : 1 - \sum_i \mathbf{x}_i \geq 0, \ \mathbf{x} \geq 0\}$ *be the n-dimensional standard simplex and let* $f \in \mathbb{R}[\mathbf{x}]$ *be a polynomial of degree d. Then for any* $r \geq 2nd$, *the lower bound* $\text{lb}(f, \mathcal{T}(\mathbf{X}))_r$ *for the minimization of* f *over* Δ^n *satisfies:*

$$f_{\min} - \text{lb}(f, \mathcal{T}(\mathbf{X}))_r \leq \frac{C_\Delta(n,d)}{r^2} \cdot (f_{\max} - f_{\min}).$$

Here, $C_\Delta(n,d)$ *is a constant depending only on* n, d. *This constant depends polynomially on n (for fixed d) and polynomially on d (for fixed n). See relation* (8.29) *for details.*

As we have seen in Section 6.2, we may obtain a convergence rate in $O(1/r^2)$ for the *upper* bounds $\text{ub}(f, \mathcal{T}(\mathbf{X}), \mu_{\mathbf{X}})_r$ on the unit ball and simplex essentially as a side result of our main proof technique. The reference measures μ_B and μ_Δ are defined below in (8.16) and (8.18). However, the obtained rates do not improve meaningfully upon previous results. Indeed, we showed in Chapter 3 that even the weaker bounds $\text{ub}(f, \mathbf{X})_r$ already converge to the global minimum at a rate in $O(1/r^2)$ for these sets (although for different reference measures).

THEOREM 8.3. *Let* $\mathbf{X} = B^n$ *be the n-dimensional unit ball equipped with the measure* μ_B *defined in* (8.16). *Let* $f \in \mathbb{R}[\mathbf{x}]$ *be a polynomial of degree d. Then for any* $r \geq d$, *the upper bound* $\text{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r$ *for the minimization of* f *over* B^n *satisfies:*

$$\text{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r - f_{\min} \leq \frac{C_B(n,d)}{2r^2} \cdot (f_{\max} - f_{\min}).$$

Here, $C_B(n,d)$ *is the constant of Theorem 8.1.*

THEOREM 8.4. *Let* $\mathbf{X} = \Delta^n$ *be the* n*-dimensional standard simplex equipped with the measure* μ_Δ *defined in* (8.18). *Let* $f \in \mathbb{R}[\mathbf{x}]$ *be a polynomial of degree* d. *Then for any* $r \geq d$, *the Lasserre-type upper bound* $\mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r$ *for the minimization of* f *over* Δ^n *satisfies:*

$$\mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu_\Delta)_r - f_{\min} \leq \frac{C_\Delta(n,d)}{2r^2} \cdot (f_{\max} - f_{\min}).$$

Here, $C_\Delta(n,d)$ *is the constant of Theorem 8.2.*

Outline of the proof. We outline how the polynomial kernel method of Chapter 6 specializes to this setting. Let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree $d \in \mathbb{N}$. We wish to show that

$$f - f_{\min} + \varepsilon \in \mathcal{T}(\mathbf{X})_{2r}$$

for some small $\varepsilon > 0$. Up to translation and scaling, we may assume that $f_{\min} = 0$ and that $\|f\|_{\mathbf{X}} := \max_{\mathbf{x} \in \mathbf{X}} |f(x)| = 1$. Let $\varepsilon > 0$. Recall that we wish to construct an (invertible) linear operator $\mathbf{K} : \mathbb{R}[\mathbf{x}]_d \to \mathbb{R}[\mathbf{x}]_d$ which satisfies the following three properties:

$$\mathbf{K}(1) = 1, \tag{P1}$$

$$\mathbf{K}p \in \mathcal{T}(\mathbf{X})_{2r} \quad \text{for all } p \in \mathcal{P}_+(\mathbf{X})_d \tag{P2}$$

$$\max_{\mathbf{x} \in \mathbf{X}} |\mathbf{K}^{-1} f(\mathbf{x}) - f(\mathbf{x})| \leq \varepsilon. \tag{P3}$$

As we saw in Section 6.1, the existence of such an operator implies that $f + \varepsilon \in \mathcal{T}(\mathbf{X})_{2r}$. Indeed, since f is nonnegative on $\mathbf{X}$ by assumption, we know that $f(\mathbf{x}) + \varepsilon \geq \varepsilon$ for $\mathbf{x} \in \mathbf{X}$. By properties (P1) and (P3), it follows that $\mathbf{K}^{-1}(f + \varepsilon) \in \mathcal{P}_+(\mathbf{X})$. Using property (P2), we may thus conclude that:

$$f + \varepsilon = \mathbf{K}\big(\mathbf{K}^{-1}(f + \varepsilon)\big) \in \mathcal{T}(\mathbf{X})_{2r},$$

meaning that $f_{\min} - \mathrm{lb}(f, \mathcal{T}(\mathbf{X}))_r \leq \varepsilon$.

The statements of Theorem 8.1 and Theorem 8.2 may thus be proven by showing the existence (for each $r \in \mathbb{N}$ large enough) of an operator $\mathbf{K}$ which satisfies (P1), (P2) and (P3) with $\varepsilon = O(1/r^2)$. We summarize this observation in the following Lemma for future reference.

LEMMA 8.5. *Let* $\mathbf{X} \subseteq \mathbb{R}^n$ *be a compact semialgebraic set and let* f *be a polynomial on* $\mathbf{X}$ *of degree* d. *Suppose that there exists a nonsingular linear operator* $\mathbf{K} : \mathbb{R}[\mathbf{x}]_d \to \mathbb{R}[\mathbf{x}]_d$ *which satisfies the properties* (P1), (P2) *and* (P3) *for certain* $\varepsilon \geq 0$. *Then* $f_{\min} - \mathrm{lb}(f, \mathcal{T}(\mathbf{X}))_r \leq \varepsilon$.

We use the polynomial kernel method to construct operators that satisfy (P1), (P2) and (P3) is: Let $\mathrm{K} : \mathbf{X} \times \mathbf{X} \to \mathbb{R}$ be a polynomial kernel on $\mathbf{X}$, meaning that $\mathrm{K}(\mathbf{x}, \mathbf{y})$ is a polynomial in the variables $\mathbf{x}, \mathbf{y}$. After choosing a measure μ supported on $\mathbf{X}$, we may associate a linear operator $\mathbf{K} : \mathbb{R}[\mathbf{x}] \to \mathbb{R}[\mathbf{x}]$

to K by setting:

$$\mathbf{K}p(\mathbf{x}) := \int_{\mathbf{X}} \mathrm{K}(\mathbf{x}, \mathbf{y}) p(\mathbf{y}) d\mu(\mathbf{y}) \quad (p \in \mathbb{R}[\mathbf{x}]). \tag{8.2}$$

By Lemma 8.8 the operator $\mathbf{K}$ satisfies (P2) if the polynomial $\mathbf{x} \mapsto \mathrm{K}(\mathbf{x}, \mathbf{y})$ lies in $\mathcal{T}(\mathbf{X})_{2r}$ for all *fixed* $\mathbf{y} \in \mathbf{X}$. Furthermore, (P1) and (P3) may be verified by analyzing the *eigenvalues* of $\mathbf{K}$; roughly speaking, $\mathbf{K}$ satisfies (P1), (P3) if its eigenvalues are sufficiently close to 1.

It remains, then, to construct a suitable kernel K on $\mathbf{X}$. For this purpose, we consider the (perturbed) Christoffel-Darboux kernel:

$$\mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda) := \sum_{k=0}^{2r} \lambda_k \mathrm{CD}^{(k)}(\mathbf{x}, \mathbf{y}) \quad (\lambda_k \in \mathbb{R}, \quad 0 \leq k \leq 2r), \tag{8.3}$$

whose associated operator has eigenvalues equal to $\lambda_0, \ldots, \lambda_{2r}$ (with multiplicity). See also (6.7). In Chapters 1 and 7, we saw that this kernel is given on the binary cube $\{0,1\}^n$ by:

$$\mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda) = \sum_{k=0}^{2r} \lambda_k \mathcal{K}_k^{(n)}(|\mathbf{x} - \mathbf{y}|_1). \tag{8.4}$$

The idea there was to select a univariate polynomial $q \in \mathbb{R}[\mathbf{x}]_r$ of degree r and consider the kernel:

$$\mathrm{K}(\mathbf{x}, \mathbf{y}) = q^2(|\mathbf{x} - \mathbf{y}|_1),$$

which clearly lies in $\mathcal{Q}(\{0,1\}^n)$. Using the closed form expression (8.4), the eigenvalues of the operator $\mathbf{K}$ associated to K are given by the coefficients λ_k in the expansion:

$$q^2(x) = \sum_{k=0}^{2r} \lambda_k \mathcal{K}_k^{(n)}(|\mathbf{x} - \mathbf{y}|_1).$$

Therefore, the analysis could be concluded by finding a q for which these coefficients are sufficiently close to 1.

The closed form expression of the Christoffel-Darboux kernel on the binary cube follows from a classical *summation formula* for the Krawtchouk polynomials, see Section 1.1 and also Section 7.1. As we saw in Chapter 1, a similar summation formula is available for Gegenbauer polynomials, which yield a closed form of the Christoffel-Darboux kernel on the hypersphere S^{n-1}. Fang and Fawzi [**FF21**] in fact used this closed form to analyze the lower bounds on S^{n-1}.

In this chapter, we shall use different summation formulas for Gegenbauer polynomials due to Xu [**Xu99, Xu98**], which yield closed form expressions of the Christoffel-Darboux kernel on the unit ball and on the standard simplex. These expressions are significantly more complicated than the ones for $\{0,1\}^n$ and S^{n-1}, but they are nonetheless very useful.

Organization. The rest of this chapter is organized as follows. In Section 8.1, we introduce some notations and cover the necessary preliminaries on orthogonal (Gegenbauer) polynomials. In Section 8.2, we present closed form expressions of the Christoffel-Darboux kernel and use them to obtain kernels whose associated operators satisfy (P2) and whose eigenvalues are given by the coefficients of a univariate sum of squares in an appropriate basis of orthogonal polynomials. In Section 8.3, we show how to choose this sum of squares so that (P1), (P3) are satisfied and finish the proof of Theorem 8.1 and Theorem 8.2. In Section 8.4, we extend our proof technique to the upper bounds to obtain Theorem 8.3 and Theorem 8.4.

8.1. Preliminaries

8.1.1. Notations. We write $\mathbb{R}[x]$ for the univariate polynomial ring, while reserving the bold-face notation $\mathbb{R}[\mathbf{x}] = \mathbb{R}[\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n]$ to denote the ring of polynomials in n variables. We denote $\|p\|_{\mathbf{X}} := \max_{\mathbf{x}\in\mathbf{X}} |p(\mathbf{x})|$ for the supremum-norm of $p \in \mathbb{R}[\mathbf{x}]$ on $\mathbf{X}$. We call a univariate polynomial p *even* if $p(x) = p(-x)$ for all $x \in \mathbb{R}$, and *odd* if $p(x) = -p(-x)$ for all $x \in \mathbb{R}$. Finally, we write $|\mathbf{x}| := \sum_{i=1}^{n} \mathbf{x}_i$ for $\mathbf{x} \in \mathbb{R}^n_{\geq 0}$.

8.1.2. Gegenbauer polynomials. We recall some properties of the *Gegenbauer polynomials* (also known as *ultraspherical* polynomials), introduced in Chapter 1. For $n \geq 2$, let $w_n(x) := c_n(1-x^2)^{\frac{n-2}{2}}$, where $c_n > 0$ is chosen so that:

$$\int_{-1}^{1} w_n(x)dx = 1.$$

The Gegenbauer polynomials $\{\mathcal{G}_k^{(\frac{n-2}{2})} : k \in \mathbb{N}\}$ are defined as the set of orthogonal polynomials on $[-1,1]$ w.r.t. the weight function w_n. That is, the polynomial $\mathcal{G}_k^{(\frac{n-2}{2})}$ is of exact degree k for each $k \in \mathbb{N}$, and the following orthogonality relations hold:

$$\int_{-1}^{1} \mathcal{G}_k^{(\frac{n-2}{2})}(x)\mathcal{G}_{k'}^{(\frac{n-2}{2})}(x)w_n(x)dx = 0 \quad (k \neq k'). \tag{8.5}$$

For notational convenience, we adopt the normalization in this chapter for which:

$$\mathcal{G}_k^{(\frac{n-2}{2})}(x) = \widetilde{\mathcal{G}}_k^{(\frac{n-2}{2})}(1)\widetilde{\mathcal{G}}_k^{(\frac{n-2}{2})}(x) \quad (x \in \mathbb{R}), \tag{8.6}$$

where $\widetilde{\mathcal{G}}_k^{(\frac{n-2}{2})}$ is the ortho*normal* Gegenbauer polynomial of degree k, i.e., satisfying:

$$\int_{-1}^{1} \widetilde{\mathcal{G}}_k^{(\frac{n-2}{2})}(x)\widetilde{\mathcal{G}}_{k'}^{(\frac{n-2}{2})}(x)w_n(x)dx = \delta_{kk'} \quad (k, k' \in \mathbb{N}).$$

We have that $\max_{x\in[-1,1]} |\mathcal{G}_k^{(\frac{n-2}{2})}(x)| = \mathcal{G}_k^{(\frac{n-2}{2})}(1)$, and so it will also be convenient to write:

$$\overline{\mathcal{G}}_k^{(\frac{n-2}{2})}(x) := \frac{\mathcal{G}_k^{(\frac{n-2}{2})}(x)}{\mathcal{G}_k^{(\frac{n-2}{2})}(1)} = \frac{\widetilde{\mathcal{G}}_k^{(\frac{n-2}{2})}(x)}{\widetilde{\mathcal{G}}_k^{(\frac{n-2}{2})}(1)} \quad (x \in \mathbb{R})$$

for the normalization of the Gegenbauer polynomials which satisfies:

$$\max_{x\in[-1,1]} |\overline{\mathcal{G}}_k^{(\frac{n-2}{2})}(x)| = \overline{\mathcal{G}}_k^{(\frac{n-2}{2})}(1) = 1 \quad (k \in \mathbb{N}).$$

The upshot is that for these choices of normalization, we have:

$$\int_{-1}^{1} \mathcal{G}_k^{(\frac{n-2}{2})}(x)\overline{\mathcal{G}}_{k'}^{(\frac{n-2}{2})}(x)w_n(x)dx = \int_{-1}^{1} \widetilde{\mathcal{G}}_k^{(\frac{n-2}{2})}(x)\widetilde{\mathcal{G}}_{k'}^{(\frac{n-2}{2})}(x)w_n(x)dx = \delta_{kk'}. \tag{8.7}$$

We will make use of the expansion:

$$q(x) = \sum_{k=0}^{d} \lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x) \quad (x \in \mathbb{R}) \tag{8.8}$$

of a univariate polynomial q of degree d in the basis of Gegenbauer polynomials. Using (8.7), the coefficients λ_k in (8.8) are given by:

$$\lambda_k = \int_{-1}^{1} \overline{\mathcal{G}}_k^{(\frac{n-2}{2})}(x)q(x)w_n(x)dx \quad (0 \le k \le d). \tag{8.9}$$

Using (8.9) and the fact that $\overline{\mathcal{G}}_k^{(\frac{n-2}{2})}(x) \le \overline{\mathcal{G}}_0^{(\frac{n-2}{2})}(x) = 1$ for every $x \in [-1,1]$ and $k \in \mathbb{N}$, we find that:

$$\lambda_k \le \lambda_0 \quad (1 \le k \le d) \tag{8.10}$$

whenever q is nonnegative on $[-1,1]$. Furthermore, we note that if q is an *even* polynomial of degree $2d$, we may write:

$$q(x) = \sum_{k=0}^{d} \lambda_{2k} \mathcal{G}_{2k}^{(\frac{n-2}{2})}(x) \quad (x \in \mathbb{R}). \tag{8.11}$$

Indeed, as the odd degree Gegenbauer polynomials are odd functions, the integral (8.9) vanishes for odd k in this case.

8.2. Construction of the linear operator

In this section, we explain how to construct a suitable linear operator $\mathbf{K}$. We recall briefly the setup of kernel operators. Let $\mathbf{X} \subseteq \mathbb{R}^n$ be a compact set and let μ be a measure whose support is exactly $\mathbf{X}$. We may define an inner product $\langle \cdot, \cdot \rangle_\mu$ on the space $\mathcal{P}(\mathbf{X})$ of polynomials on $\mathbf{X}$ by setting:

$$\langle p, q \rangle_\mu := \int_{\mathbf{X}} p(\mathbf{x})q(\mathbf{x})d\mu(\mathbf{x}) \quad (p, q \in \mathcal{P}(\mathbf{X})).$$

We write $\{P_\alpha : \alpha \in \mathbb{N}^n\}$ for an orthonormal basis of $\mathcal{P}(\mathbf{X})$ w.r.t. this inner product, ordered so that P_α is of exact degree $|\alpha| = \sum_{i=1}^n \alpha_i$ for all $\alpha \in \mathbb{N}^n$. The Christoffel-Darboux kernel $\mathrm{CD}_r : \mathbf{X} \times \mathbf{X} \to \mathbb{R}$ of degree $r \in \mathbb{N}$ for $(\mathbf{X}, \mu)$ is then defined as:

$$\mathrm{CD}_r(\mathbf{x}, \mathbf{y}) = \sum_{\alpha \in \mathbb{N}^n_r} P_\alpha(\mathbf{x}) P_\alpha(\mathbf{y}) \quad (\mathbf{x}, \mathbf{y} \in \mathbf{X}). \tag{8.12}$$

From the orthonormality of the P_α, it follows that the operator $\mathbf{CD}_r$ associated to the Christoffel-Darboux kernel CD_r via (6.6) acts as the identity on $\mathcal{P}(\mathbf{X})_r$. That is, we have:

$$\mathbf{CD}_r p(\mathbf{x}) = \int_{\mathbf{X}} \mathrm{CD}_r(\mathbf{x}, \mathbf{y}) p(\mathbf{y}) d\mu(\mathbf{y}) = p(\mathbf{x}) \quad (\mathbf{x} \in \mathbf{X},\ p \in \mathcal{P}(\mathbf{X})_r).$$

For $k \in \mathbb{N}$, we write $H_k := \mathrm{span}\{P_\alpha : |\alpha| = k\}$ for the subspace of $\mathcal{P}(\mathbf{X})$ spanned by the P_α of exact degree k. Note that we may equivalently define H_k as:

$$H_k := \{p \in \mathcal{P}(\mathbf{X})_k : \langle p, q \rangle_\mu = 0 \text{ for all } q \in \mathcal{P}(\mathbf{X})_{k-1}\}.$$

In particular, we see that H_k does not depend on our choice of basis $\{P_\alpha\}$, but only on the measure μ. In light of this fact, it is convenient to adopt the vector-notation:

$$\mathbb{P}_k(\mathbf{x}) := (P_\alpha(\mathbf{x}))_{|\alpha|=k} \quad (k \in \mathbb{N}).$$

The kernel $\mathrm{CD}^{(k)}(\mathbf{x}, \mathbf{y}) := \mathbb{P}_k(\mathbf{x})^\top \mathbb{P}_k(\mathbf{y})$ does not depend on the choice of basis $\{P_\alpha\}$, and its associated operator reproduces the subspace H_k. That is, if we decompose a polynomial $p \in \mathcal{P}(\mathbf{X})$ as:

$$p(\mathbf{x}) = \sum_{k=0}^{\deg(p)} p_k(\mathbf{x}) \quad (p_k \in H_k), \tag{8.13}$$

we then have that:

$$\mathbf{CD}^{(k)} p(\mathbf{x}) = \int_{\mathbf{X}} \mathbb{P}_k(\mathbf{x})^\top \mathbb{P}_k(\mathbf{y}) p(\mathbf{y}) d\mu(\mathbf{y}) = p_k(\mathbf{x}) \quad (\mathbf{x} \in \mathbf{X},\ k \in \mathbb{N}). \tag{8.14}$$

After regrouping the terms in (8.12), we can express the Christoffel-Darboux kernel as:

$$\mathrm{CD}_r(\mathbf{x}, \mathbf{y}) = \sum_{k=0}^{r} \mathrm{CD}^{(k)}(\mathbf{x}, \mathbf{y}) = \sum_{k=0}^{r} \mathbb{P}_k(\mathbf{x})^\top \mathbb{P}_k(\mathbf{y}). \tag{8.15}$$

As we have seen, a closed form expression for CD_r may be derived based on the regrouping (8.15) on the binary cube and the unit sphere. See relations (1.15) and (1.17). In these special cases, the term $\mathbb{P}_k(\mathbf{x})^\top \mathbb{P}_k(\mathbf{y})$ may be expressed by composing a *univariate* polynomial with a relatively simple multivariate polynomial.

As we see now, such expressions also exist for the unit ball and the simplex.

8.2.1. A closed form for the unit ball. Consider the unit ball $B^n := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|^2 \leq 1\} \subseteq \mathbb{R}^n$, which we equip with the $O(n)$-invariant probability measure μ_B given by:

$$d\mu_B(\mathbf{x}) = c_n(1 - \|\mathbf{x}\|^2)^{-\frac{1}{2}} d\mathbf{x} \quad (\mathbf{x} \in B^n), \tag{8.16}$$

with $c_n > 0$ a normalization constant. As B^n is full-dimensional, we have $\mathcal{P}(B^n) = \mathbb{R}[\mathbf{x}]$. Xu derives the following closed form of the Christoffel-Darboux kernel on B^n.

Theorem 8.6 (Xu [**Xu99**], Theorem 3.1). *Let $\{P_\alpha : \alpha \in \mathbb{N}\}$ be an orthonormal basis of $\mathcal{P}(B^n)$ w.r.t. μ_B. Then the P_α satisfy the following summation formula in terms of the Gegenbauer polynomials*[1] (8.5)*:*

$$\begin{aligned}\mathbb{P}_k(\mathbf{x})^\top \mathbb{P}_k(\mathbf{y}) = \frac{1}{2} \cdot \Big(&\mathcal{G}_k^{(\frac{n-2}{2})}(\mathbf{x}\cdot\mathbf{y} + \sqrt{1-\|\mathbf{x}\|^2}\sqrt{1-\|\mathbf{y}\|^2}) \quad + \\ &\mathcal{G}_k^{(\frac{n-2}{2})}(\mathbf{x}\cdot\mathbf{y} - \sqrt{1-\|\mathbf{x}\|^2}\sqrt{1-\|\mathbf{y}\|^2})\Big)\end{aligned} \quad (\mathbf{x}, \mathbf{y} \in B^n). \tag{8.17}$$

Using (8.17)*, we have the following closed form of the Christoffel-Darboux kernel* CD_r*:*

$$\begin{aligned}\mathrm{CD}_r(\mathbf{x}, \mathbf{y}) = \frac{1}{2}\sum_{k=0}^{r} \Big(&\mathcal{G}_k^{(\frac{n-2}{2})}(\mathbf{x}\cdot\mathbf{y} + \sqrt{1-\|\mathbf{x}\|^2}\sqrt{1-\|\mathbf{y}\|^2}) \quad + \\ &\mathcal{G}_k^{(\frac{n-2}{2})}(\mathbf{x}\cdot\mathbf{y} - \sqrt{1-\|\mathbf{x}\|^2}\sqrt{1-\|\mathbf{y}\|^2})\Big)\end{aligned} \quad (\mathbf{x}, \mathbf{y} \in B^n).$$

8.2.2. A closed form for the standard simplex. Consider the standard simplex:

$$\Delta^n := \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x} \geq 0,\ 1 - |\mathbf{x}| \geq 0\} \subseteq \mathbb{R}^n.$$

We equip Δ^n with the probability measure μ_Δ given by:

$$d\mu_\Delta(\mathbf{x}) = c_n \mathbf{x}_1^{-1/2}\mathbf{x}_2^{-1/2} \ldots \mathbf{x}_n^{-1/2}(1 - |\mathbf{x}|)^{-1/2} d\mathbf{x} \quad (\mathbf{x} \in \Delta^n), \tag{8.18}$$

where $c_n > 0$ is a normalization constant. As the simplex is full-dimensional, we have $\mathcal{P}(\Delta^n) = \mathbb{R}[\mathbf{x}]$. Xu derives the following closed form of the Christoffel-Darboux kernel.

Theorem 8.7 (Xu [**Xu98**], Corollary 2.4). *Let $\{P_\alpha : \alpha \in \mathbb{N}\}$ be an orthonormal basis of $\mathcal{P}(\Delta^n)$ w.r.t. μ_Δ. Then the P_α satisfy the following summation formula in terms of the Gegenbauer polynomials* (8.5)*:*

$$\mathbb{P}_k(\mathbf{x})^\top \mathbb{P}_k(\mathbf{y}) = \frac{1}{2^{n+1}} \sum_{t \in \{-1,1\}^{n+1}} \mathcal{G}_{2k}^{(\frac{n-2}{2})}\Big(\sum_{i=1}^{n+1} \sqrt{\mathbf{x}_i \mathbf{y}_i} t_i\Big) \quad (\mathbf{x}, \mathbf{y} \in \Delta^n). \tag{8.19}$$

---

[1] Note that the Gegenbauer polynomial of degree k in [**Xu99**] differs by a factor $(k + \frac{n-1}{2})/\frac{n-1}{2}$ from the one used here (compare (2.10) in [**Xu99**] to (8.6)).

Here, we write $\mathbf{x}_{n+1} := 1 - |\mathbf{x}|$, $\mathbf{y}_{n+1} := 1 - |\mathbf{y}|$. *Using* (8.19), *we have the following closed form of the Christoffel-Darboux kernel* CD_r*:*

$$\mathrm{CD}_r(\mathbf{x},\mathbf{y}) = \frac{1}{2^{n+1}} \sum_{k=0}^{r} \sum_{t\in\{-1,1\}^{n+1}} \mathcal{G}_{2k}^{(\frac{n-2}{2})}\Big(\sum_{i=1}^{n+1} \sqrt{\mathbf{x}_i\mathbf{y}_i} t_i\Big) \quad (\mathbf{x},\mathbf{y} \in \Delta^n).$$

8.2.3. Sum-of-squares representations. Based on the closed forms of the Christoffel-Darboux kernel derived above, we may define kernels $\mathrm{K}(\mathbf{x},\mathbf{y})$ for the unit ball and simplex whose associated operators satisfy property (P2). Recall from Lemma 6.2; that it suffices for (P2) to hold that $\mathbf{x} \mapsto \mathrm{K}(\mathbf{x},\mathbf{y})$ lies $\mathcal{T}(\mathbf{X})_{2r}$ for all $\mathbf{y} \in \mathbf{X}$ fixed.

Lemma 8.8 (Restatement of Lemma 6.2). *Let* $\mathbf{X} \subseteq \mathbb{R}^n$ *be a compact semi-algebraic set, and let* μ *be a finite measure supported on* $\mathbf{X}$. *Let* $Q \subseteq \mathbb{R}[\mathbf{x}]$ *be a convex cone, and suppose that* $\mathrm{K} : \mathbf{X} \times \mathbf{X} \to \mathbb{R}$ *is a polynomial kernel for which* $\mathrm{K}(\cdot,\mathbf{y}) \in Q$ *for each* $\mathbf{y} \in \mathbf{X}$ *fixed. Then if* $p \in \mathbb{R}[\mathbf{x}]$ *is nonnegative on* $\mathbf{X}$, *we have* $\mathbf{K}p \in Q$. *That is, when selecting* $Q = \mathcal{T}(\mathbf{X})_{2r}$, *the operator* $\mathbf{K}$ *associated to* K *satisfies* (P2).

In Chapter 7, it was very straightforward to see that the kernel $\mathrm{K}(\mathbf{x},\mathbf{y}) = q^2(|\mathbf{x}-\mathbf{y}|_1)$ lies in $\mathcal{Q}(\{0,1\}^n)_{2r}$ for any polynomial $q \in \mathbb{R}[\mathbf{x}]$ of degree at most r. Turning now to the unit ball and simplex, the situation is slightly more complicated.

The unit ball. Let $q \in \Sigma[x]_{2r}$ be a univariate sum of squares, with expansion $q(x) = \sum_{k=0}^{2r} \lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x)$ in the basis of Gegenbauer polynomials (8.8). In light of the closed form (8.17) of the Christoffel-Darboux kernel on the unit ball, we have:

$$\begin{aligned}\mathrm{CD}_{2r}(\mathbf{x},\mathbf{y};\lambda) = \frac{1}{2}\Big(q(\mathbf{x}\cdot\mathbf{y} + \sqrt{1-\|\mathbf{x}\|^2}\sqrt{1-\|\mathbf{y}\|^2}) \quad + \\ q(\mathbf{x}\cdot\mathbf{y} - \sqrt{1-\|\mathbf{x}\|^2}\sqrt{1-\|\mathbf{y}\|^2})\Big)\end{aligned} \qquad (\mathbf{x},\mathbf{y} \in B^n). \tag{8.20}$$

Lemma 8.9. *Let* $q \in \Sigma[x]_{2r}$ *be a univariate sum of squares. Then the kernel* $\mathrm{CD}_{2r}(\mathbf{x},\mathbf{y};\lambda)$ *in* (8.20) *satisfies* $\mathrm{CD}_{2r}(\cdot,\mathbf{y};\lambda) \in \mathcal{T}(B^n)_{2r}$ *for fixed* $\mathbf{y} \in B^n$. *As a result, its associated operator satisfies* (P2) *by Lemma 8.8.*

The simplex. Let $q(x) = \sum_{k=0}^{4r} \lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x)$ again be a univariate sum of squares, now of degree $4r$. Using (8.11), we have:

$$q_{\mathrm{even}}(x) := \frac{q(x)+q(-x)}{2} = \sum_{k=0}^{2r} \lambda_{2k} \mathcal{G}_{2k}^{(\frac{n-2}{2})}(x) \quad (x \in \mathbb{R}). \tag{8.21}$$

In light of the closed form (8.19) of the Christoffel-Darboux kernel on the simplex, we find:

$$\mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda_{\mathrm{even}}) = \frac{1}{2^{n+1}} \sum_{t \in \{-1,1\}^{n+1}} q_{\mathrm{even}}\Big(\sum_{i=1}^{n+1} \sqrt{\mathbf{x}_i \mathbf{y}_i} t_i\Big) \quad (\mathbf{x}, \mathbf{y} \in \Delta^n). \quad (8.22)$$

Here, $\lambda_{\mathrm{even}} := (\lambda_{2k})_{0 \leq k \leq 2r}$ and $\mathbf{x}_{n+1} = 1 - |\mathbf{x}|$, $\mathbf{y}_{n+1} = 1 - |\mathbf{y}|$.

LEMMA 8.10. *Let $q \in \Sigma[x]_{4r}$ be a univariate sum of squares of degree $4r$, and let q_{even} be as in* (8.21)*. Then the kernel* $\mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda_{\mathrm{even}})$ *in* (8.22) *satisfies* $\mathrm{CD}_{2r}(\cdot, \mathbf{y}; \lambda_{\mathrm{even}}) \in \mathcal{T}(\Delta^n)_{2r}$ *for fixed $\mathbf{y} \in \Delta^n$. As a result, its associated operator satisfies* (P2) *by Lemma 8.8.*

For the proof of Lemma 8.9 and Lemma 8.10, we need the following lemma.

LEMMA 8.11. *Let $p \in \mathbb{R}[x]$ be a univariate polynomial of degree r. Let $\mathbf{u}, \mathbf{v}$ be formal variables. Then the polynomial $p(\mathbf{u}+\mathbf{v})^2 + p(\mathbf{u}-\mathbf{v})^2$ admits a representation:*

$$p(\mathbf{u}+\mathbf{v})^2 + p(\mathbf{u}-\mathbf{v})^2 = \mathbf{v}^2 h_{\mathrm{odd}}(\mathbf{u}, \mathbf{v}^2)^2 + h_{\mathrm{even}}(\mathbf{u}, \mathbf{v}^2)^2,$$

where $\mathbf{v} h_{\mathrm{odd}}(\mathbf{u}, \mathbf{v}^2)$ and $h_{\mathrm{even}}(\mathbf{u}, \mathbf{v}^2) \in \mathbb{R}[\mathbf{u}, \mathbf{v}]$ are polynomials of degree r.

PROOF. For convenience, let $u \in \mathbb{R}[\mathbf{u}, \mathbf{v}]$ be given by $u(\mathbf{u}, \mathbf{v}) = \mathbf{u} + \mathbf{v}$, so that:

$$p(\mathbf{u}+\mathbf{v})^2 + p(\mathbf{u}-\mathbf{v})^2 = p(u(\mathbf{u}, \mathbf{v}))^2 + p(u(\mathbf{u}, -\mathbf{v}))^2 = h(\mathbf{u}, \mathbf{v})^2 + h(\mathbf{u}, -\mathbf{v})^2,$$

where $h = p \circ u \in \mathbb{R}[\mathbf{u}, \mathbf{v}]_r$. If we expand h in the monomial basis of $\mathbb{R}[\mathbf{u}, \mathbf{v}]$ as:

$$h(\mathbf{u}, \mathbf{v}) = \sum_{i+j \leq r} h_{ij} \mathbf{u}^i \mathbf{v}^j \quad (h_{ij} \in \mathbb{R}),$$

we may perform the following computation (where all summations are taken over $i+j \leq r$):

$$\begin{aligned}
&h(\mathbf{u},\mathbf{v})^2 + h(\mathbf{u},-\mathbf{v})^2 \\
&\quad= \Big(\sum h_{ij}\mathbf{u}^i\mathbf{v}^j\Big)^2 + \Big(\sum h_{ij}\mathbf{u}^i(-\mathbf{v})^j\Big)^2 \\
&\quad= \Big(\sum_{j\text{ odd}} h_{ij}\mathbf{u}^i\mathbf{v}^j + \sum_{j\text{ even}} h_{ij}\mathbf{u}^i\mathbf{v}^j\Big)^2 + \Big(\sum_{j\text{ odd}} h_{ij}\mathbf{u}^i(-\mathbf{v})^j + \sum_{j\text{ even}} h_{ij}\mathbf{u}^i(-\mathbf{v})^j\Big)^2 \\
&\quad= \Big(\sum_{j\text{ odd}} h_{ij}\mathbf{u}^i\mathbf{v}^j + \sum_{j\text{ even}} h_{ij}\mathbf{u}^i\mathbf{v}^j\Big)^2 + \Big(-\sum_{j\text{ odd}} h_{ij}\mathbf{u}^i\mathbf{v}^j + \sum_{j\text{ even}} h_{ij}\mathbf{u}^i\mathbf{v}^j\Big)^2 \\
&\quad= (\sum_{j\text{ odd}} h_{ij}\mathbf{u}^i\mathbf{v}^j)^2 + (\sum_{j\text{ even}} h_{ij}\mathbf{u}^i\mathbf{v}^j)^2 + 2(\sum_{j\text{ odd}} h_{ij}\mathbf{u}^i\mathbf{v}^j)(\sum_{j\text{ even}} h_{ij}\mathbf{u}^i\mathbf{v}^j) \\
&\qquad+ (\sum_{j\text{ odd}} h_{ij}\mathbf{u}^i\mathbf{v}^j)^2 + (\sum_{j\text{ even}} h_{ij}\mathbf{u}^i\mathbf{v}^j)^2 - 2(\sum_{j\text{ odd}} h_{ij}\mathbf{u}^i\mathbf{v}^j)(\sum_{j\text{ even}} h_{ij}\mathbf{u}^i\mathbf{v}^j) \\
&\quad= 2(\sum_{j\text{ odd}} h_{ij}\mathbf{u}^i\mathbf{v}^j)^2 + 2(\sum_{j\text{ even}} h_{ij}\mathbf{u}^i\mathbf{v}^j)^2 \\
&\quad= 2\mathbf{v}^2(\sum_{j\text{ odd}} h_{ij}\mathbf{u}^i\mathbf{v}^{j-1})^2 + 2(\sum_{j\text{ even}} h_{ij}\mathbf{u}^i\mathbf{v}^j)^2.
\end{aligned}$$

But now we see that there exist $h_{\text{odd}}, h_{\text{even}} \in \mathbb{R}[\mathbf{u},\mathbf{v}]$ of appropriate degree such that:

$$p(\mathbf{u}+\mathbf{v})^2 + p(\mathbf{u}-\mathbf{v})^2 = h(\mathbf{u},\mathbf{v})^2 + h(\mathbf{u},-\mathbf{v})^2 = \mathbf{v}^2 h_{\text{odd}}(\mathbf{u},\mathbf{v}^2)^2 + h_{\text{even}}(\mathbf{u},\mathbf{v}^2)^2,$$

as required. □

Proof of Lemma 8.9. We may assume w.l.o.g. that $q = p^2$ is a square. For $\mathbf{x},\mathbf{y} \in B^n$, write $\mathbf{u} = \mathbf{x}\cdot\mathbf{y}$ and $\mathbf{v} = \sqrt{1-\|\mathbf{x}\|^2}\sqrt{1-\|\mathbf{y}\|^2}$, so that:

$$\mathrm{CD}_{2r}(\mathbf{x},\mathbf{y};\lambda) = p(\mathbf{u}+\mathbf{v})^2 + p(\mathbf{u}-\mathbf{v})^2.$$

By Lemma 8.11, there exist $h_{\text{odd}}, h_{\text{even}} \in \mathbb{R}[\mathbf{v},\mathbf{u}]$ of appropriate degree so that:

$$\mathrm{CD}_{2r}(\mathbf{x},\mathbf{y};\lambda) = \mathbf{v}^2 h_{\text{odd}}(\mathbf{u},\mathbf{v}^2)^2 + h_{\text{even}}(\mathbf{u},\mathbf{v}^2)^2,$$

which lies in $\mathcal{T}(B^n)_{2r}$ for $\mathbf{y} \in B^n$ as $\mathbf{u} = \mathbf{x}\cdot\mathbf{y}$ and $\mathbf{v}^2 = (1-\|\mathbf{x}\|^2)(1-\|\mathbf{y}\|^2)$. □

Proof of Lemma 8.10. Note first that $q_{\text{even}}(x) = \frac{1}{2}\big(q(x)+q(-x)\big)$ is itself a sum of squares. In view of (8.22), it now suffices to show that for any square p^2 of degree $4r$, the kernel:

$$\mathrm{K}(\mathbf{x},\mathbf{y}) := \sum_{t\in\{-1,1\}^{n+1}} p^2\Big(\sum_{i=1}^{n+1}\sqrt{\mathbf{x}_i\mathbf{y}_i}\,t_i\Big)$$

lies in $\mathcal{T}(\Delta^n)_{2r}$ for fixed $\mathbf{y}\in\Delta^n$ (recall that we write $\mathbf{x}_{n+1} = 1-|\mathbf{x}|, \mathbf{y}_{n+1} = 1-|\mathbf{y}|$). We prove this by successive application of Lemma 8.11. Set first

$\mathbf{u}_{(0)} = \sum_{i=1}^{n} \sqrt{\mathbf{x}_i \mathbf{y}_i} t_i$ and $\mathbf{v}_{(0)} = \sqrt{\mathbf{x}_{n+1}\mathbf{y}_{n+1}}$. Then Lemma 8.11 tell us that:

$$\begin{aligned}
\mathrm{K}(\mathbf{x}, \mathbf{y}) &= \sum_{t \in \{-1,1\}^{n+1}} p^2(\mathbf{u}_{(0)} + \mathbf{v}_{(0)} t_{n+1}) \\
&= \sum_{t \in \{-1,1\}^{n}} \big(p^2(\mathbf{u}_{(0)} + \mathbf{v}_{(0)}) + p^2(\mathbf{u}_{(0)} - \mathbf{v}_{(0)})\big) \\
&= \sum_{t \in \{-1,1\}^{n}} \big(\mathbf{v}_{(0)}^2 h_0(\mathbf{u}_{(0)}, \mathbf{v}_{(0)}^2)^2 + h_1(\mathbf{u}_{(0)}, \mathbf{v}_{(0)}^2)^2\big),
\end{aligned}$$

for polynomials h_0, h_1 of appropriate degree. We may set $\mathbf{u}_{(1)} = \sum_{i=1}^{n-1} \sqrt{\mathbf{x}_i \mathbf{y}_i} t_i$ and $\mathbf{v}_{(1)} = \sqrt{\mathbf{x}_n \mathbf{y}_n}$ and proceed to find:

$$\begin{aligned}
\mathrm{K}(\mathbf{x}, \mathbf{y}) = &\sum_{t \in \{-1,1\}^{n}} \big(\mathbf{v}_{(0)}^2 h_0(\mathbf{u}_{(0)}, \mathbf{v}_{(0)}^2)^2 + h_1(\mathbf{u}_{(0)}, \mathbf{v}_{(0)}^2)^2\big) \\
= &\sum_{t \in \{-1,1\}^{n-1}} \Big(\mathbf{v}_{(0)}^2 \big(h_0^2(\mathbf{u}_{(1)} + \mathbf{v}_{(1)}, \mathbf{v}_{(0)}^2) + h_0^2(\mathbf{u}_{(1)} - \mathbf{v}_{(1)}, \mathbf{v}_{(0)}^2)\big) \\
&\qquad + h_1^2(\mathbf{u}_{(1)} + \mathbf{v}_{(1)}, \mathbf{v}_{(0)}^2) + h_1^2(\mathbf{u}_{(1)} - \mathbf{v}_{(1)}, \mathbf{v}_{(0)}^2)\Big) \\
= &\sum_{t \in \{-1,1\}^{n-1}} \Big(\mathbf{v}_{(0)}^2 \mathbf{v}_{(1)}^2 h_{00}^2(\mathbf{u}_{(1)}, \mathbf{v}_{(1)}^2, \mathbf{v}_{(0)}^2) + \mathbf{v}_{(0)}^2 h_{01}^2(\mathbf{u}_{(1)}, \mathbf{v}_{(1)}^2, \mathbf{v}_{(0)}^2) \\
&\qquad + \mathbf{v}_{(1)}^2 h_{10}^2(\mathbf{u}_{(1)}, \mathbf{v}_{(1)}^2, \mathbf{v}_{(0)}^2) + h_{11}^2(\mathbf{u}_{(1)}, \mathbf{v}_{(1)}^2, \mathbf{v}_{(0)}^2)\Big).
\end{aligned}$$

for polynomials $h_{00}, h_{01}, h_{10}, h_{11}$ of appropriate degree. After n applications of this procedure, we find that:

$$\mathrm{K}(\mathbf{x}, \mathbf{y}) = \sum_{a \in \{0,1\}^{n+1}} \prod_{i=0}^{n} \mathbf{v}_{(i)}^{2a_i} \cdot h_a^2(\mathbf{v}_{(0)}^2, \ldots, \mathbf{v}_{(n)}^2), \tag{8.23}$$

for polynomials h_a of appropriate degree, where $\mathbf{v}_{(i)}^2 = \mathbf{x}_{n+1-i}\mathbf{y}_{n+1-i}$ for $0 \leq i \leq n$. As $\Delta^n = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{x}_i \geq 0 \quad (1 \leq i \leq n+1)\}$, this means that $\mathrm{K}(\mathbf{x}, \mathbf{y}) \in \mathcal{T}(\Delta^n)_{2r}$ for $\mathbf{y} \in \Delta^n$ fixed. Take note that while the summands in (8.23) are of degree $4r$ in the variables $\mathbf{v}_{(1)}, \ldots, \mathbf{v}_{(n)}$, they indeed have degree $2r$ in the variables $\mathbf{x}_1, \ldots, \mathbf{x}_n$. □

8.3. Analysis of the linear operator

For $\mathbf{X} = B^n, \Delta^n$ and $r \in \mathbb{N}$, let $\mathbf{K} = \mathbf{CD}_{2r}(\lambda) : \mathcal{P}(\mathbf{X})_{2r} \to \mathcal{P}(\mathbf{X})_{2r}$ be the operator associated to the perturbed Christoffel-Darboux kernel $\mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda)$ defined in (8.3). As we have shown in Lemma 8.9 and Lemma 8.10 above, the operator $\mathbf{K}$ satisfies (P2) for $Q = \mathcal{T}(\mathbf{X})_{2r}$ if $\lambda = (\lambda_k)$ (resp. $\lambda_{\mathrm{ev}} = (\lambda_{2k})$) is given by the coefficients of a univariate sum of squares $q \in \sigma[x]_{2r}$ in the expansion $q(x) = \sum_{k=0}^{2r} \lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x)$ in the basis of Gegenbauer polynomials (8.8).

We now expand on the second claim made in Section 8; namely that (P3) may be expressed in terms of the difference between the λ_k and 1.

Let $q(x) = \sum_{k=0}^{2} r\lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x)$ be a univariate sum of squares. In the above, we have shown that the operator $\mathbf{K}$ corresponding to the perturbed Christoffel-Darboux kernel (8.20) on the ball (resp., (8.22) on this simplex) satisfies (P2).

Recall from (8.14) that for any polynomial p on $\mathbf{X}$ of degree d, we have $\mathbf{K}p = \sum_{k=0}^{d} \lambda_k p_k$, where $p = \sum_{k=0}^{d} p_k$ is the decomposition of (8.13). We see immediately that $\mathbf{K}$ satisfies (P1) if and only if $\lambda_0 = 1$ (i.e., when $\mathbf{K}(1) = \lambda_0 = 1$). We now expand on the second claim made in Section 8; namely that (P3) may be shown for the operator $\mathbf{K}$ by analyzing the difference between the coefficients λ_k and 1.

Recall that we consider in (P3) a polynomial f on $\mathbf{X}$ of degree d, whose sup-norm $\|f\|_{\mathbf{X}}$ over $\mathbf{X}$ is at most 1 by assumption, and that we wish to bound $\|\mathbf{K}^{-1}f - f\|_{\mathbf{X}}$. Assuming that $\lambda_0 = 1$ and $\lambda_k \neq 0$ for $1 \leq k \leq d$, we have $\mathbf{K}^{-1}f = \sum_{k=0}^{d}(1 - 1/\lambda_k)f_k$ and so:

$$\|\mathbf{K}^{-1}f - f\|_{\mathbf{X}} = \|\sum_{k=1}^{d}(1 - 1/\lambda_k)f_k\|_{\mathbf{X}} \leq \max_{1\leq k\leq d} \|f_k\|_{\mathbf{X}} \cdot \sum_{k=1}^{d} |1 - 1/\lambda_k|.$$

We have shown the following.

Lemma 8.12. *Let* $\mathbf{K}$ *be the operator associated to the perturbed Christoffel-Darboux kernel* $\mathrm{K}(\mathbf{x}, \mathbf{y}) := \mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda)$ *defined in* (8.3) *of degree* $2r$ *for certain* $\lambda = (\lambda_k)_{0\leq k\leq 2r}$. *Then* $\mathbf{K}$ *satisfies* (P1) *if* $\lambda_0 = 1$, *and it satisfies* (P3) *with:*

$$\varepsilon = \max_{1\leq k\leq d} \|f_k\|_{\mathbf{X}} \cdot \sum_{k=1}^{d} |1 - 1/\lambda_k|. \tag{8.24}$$

We now work to bound the quantity (8.24).

8.3.1. The harmonic constant. In light of the factor $\max_{1\leq k\leq d} \|f_k\|_{\mathbf{X}}$ in (8.24), we define the parameter:

$$\gamma(\mathbf{X})_d := \max_{p\in\mathbb{R}[\mathbf{x}]_d} \max_{0\leq k\leq d} \frac{\|p_k\|_{\mathbf{X}}}{\|p\|_{\mathbf{X}}}. \tag{8.25}$$

for any compact semialgebraic set $\mathbf{X}$ (equipped with a measure μ). Recall that we already studied this parameter for $\mathbf{X} = \mathbb{B}^n$ in Chapter 7. Note that $\max_{1\leq k\leq d} \|f_k\|_{\mathbf{X}} \leq \gamma(\mathbf{X})_d$ by definition. Let us first remark that $\gamma(\mathbf{X})_d < \infty$. Indeed, this follows immediately from equivalence of norms on the finite-dimensional vector space $\mathcal{P}(\mathbf{X})_d$. In particular, $\gamma(B^n)_d$ and $\gamma(\Delta^n)_d$ are finite constants depending only on n and d.

For special cases of $\mathbf{X}$, more can be shown. For instance, when $\mathbf{X}$ is the hypersphere or binary cube, the constant $\gamma(\mathbf{X})_d$ may be bounded *independently* of the dimension n [**FF21, SL21b**]. For the unit ball and simplex, we have the following.

PROPOSITION 8.13. *Let $\gamma(B^n)_d$ and $\gamma(\Delta^n)_d$ be the constants of* (8.25) *on the unit ball and simplex, respectively. Then $\gamma(B^n)_d$ and $\gamma(\Delta^n)_d$ depend polynomially on n (when d is fixed) and polynomially on d (when n is fixed).*

PROOF. Let $\mathbf{X} = B^n, \Delta^n$, equipped with the probability measure $\mu_{\mathbf{X}} = \mu_B,\ \mu_\Delta$, respectively. Let $p \in \mathbb{R}[\mathbf{x}]_d$ be a polynomial of degree d and assume that $\|p\|_{\mathbf{X}} = 1$. For $k \leq d$, we know from (8.14) that:

$$p_k(\mathbf{x}) = \mathbf{CD}^{(k)}p(\mathbf{x}) = \int_{\mathbf{X}} \mathrm{CD}^{(k)}(\mathbf{x},\mathbf{y})p(\mathbf{y})d\mu_{\mathbf{X}}(\mathbf{y}) \quad (\mathbf{x} \in \mathbf{X}).$$

Using the fact that $\|p\|_{\mathbf{X}} = 1$ and $\mu_{\mathbf{X}}$ is a probability measure, as well as the Cauchy-Schwarz inequality, we find for $\mathbf{x} \in \mathbf{X}$ that:

$$\begin{aligned}|p_k(\mathbf{x})|^2 &= |\int_{\mathbf{X}} \mathrm{CD}^{(k)}(\mathbf{x},\mathbf{y})p(\mathbf{y})d\mu_{\mathbf{X}}(\mathbf{y})|^2 \\ &\leq \int_{\mathbf{X}} \mathrm{CD}^{(k)}(\mathbf{x},\mathbf{y})^2 d\mu_{\mathbf{X}}(\mathbf{y}) \cdot \int_{\mathbf{X}} p(\mathbf{y})^2 d\mu_{\mathbf{X}}(\mathbf{y}) \\ &\leq \int_{\mathbf{X}} \mathrm{CD}^{(k)}(\mathbf{x},\mathbf{y})^2 d\mu_{\mathbf{X}}(\mathbf{y})\end{aligned}$$

Using (8.14) again, we have:

$$\int_{\mathbf{X}} \mathrm{CD}^{(k)}(\mathbf{x},\mathbf{y})^2 d\mu_{\mathbf{X}}(\mathbf{y}) = \mathrm{CD}^{(k)}(\mathbf{x},\mathbf{x}) \quad (\mathbf{x} \in \mathbf{X}). \tag{8.26}$$

It follows that:

$$\gamma(\mathbf{X})_d^2 \leq \max_{0\leq k\leq d}\max_{\mathbf{x}\in\mathbf{X}} \mathrm{CD}^{(k)}(\mathbf{x},\mathbf{x}).$$

The closed forms (8.17) and (8.19) of $\mathrm{CD}^{(k)}$ allow us to bound (8.26). On the ball, we have:

$$\mathrm{CD}^{(k)}(\mathbf{x},\mathbf{x}) = \frac{1}{2}\cdot\big(\mathcal{G}_k^{(\frac{n-2}{2})}(1) + \mathcal{G}_k^{(\frac{n-2}{2})}(2\|\mathbf{x}\|^2 - 1)\big) \quad (\mathbf{x} \in B^n).$$

In particular, we have:

$$\begin{aligned}\gamma(B^n)_d^2 &\leq \max_{0\leq k\leq d}\big(\max_{\mathbf{x}\in B^n} \mathrm{CD}^{(k)}(\mathbf{x},\mathbf{x})\big) \\ &\leq \max_{0\leq k\leq d}\big(\max_{-1\leq x\leq 1} |\mathcal{G}_k^{(\frac{n-2}{2})}(x)|\big) = \max_{0\leq k\leq d} \mathcal{G}_k^{(\frac{n-2}{2})}(1).\end{aligned}$$

On the simplex, we similarly have:

$$\mathrm{CD}^{(k)}(\mathbf{x},\mathbf{x}) = \frac{1}{2^{n+1}} \sum_{t\in\{-1,1\}^{n+1}} \mathcal{G}_{2k}^{(\frac{n-2}{2})}\big(\sum_{i=1}^{n+1} \mathbf{x}_i t_i\big) \quad (\mathbf{x} \in \Delta^n)$$

and so:

$$\gamma(\Delta^n)_d^2 \leq \max_{0\leq k\leq d} \big(\max_{\mathbf{x}\in\Delta^n} \mathrm{CD}^{(k)}(\mathbf{x},\mathbf{x})\big)$$
$$\leq \max_{0\leq k\leq d} \big(\max_{-1\leq x\leq 1} |\mathcal{G}_{2k}^{(\frac{n-2}{2})}(x)|\big) = \max_{0\leq k\leq d} \mathcal{G}_{2k}^{(\frac{n-2}{2})}(1).$$

Finally, we note that (see, e.g., (2.9) in [**Xu99**]):

$$\mathcal{G}_k^{(\frac{n-2}{2})}(1) = (1 + \frac{2k}{n-1}) \cdot \binom{k+n-2}{k} \qquad (n, k \in \mathbb{N}).$$

We conclude that the constant $\gamma(B^n)_d$ satisfies:

$$\gamma(B^n)_d^2 \leq \max_{0\leq k\leq d} \mathcal{G}_k^{(\frac{n-2}{2})}(1) = \max_{0\leq k\leq d} (1 + \frac{2k}{n-1}) \cdot \binom{k+n-2}{k}.$$

The constant $\gamma(\Delta^n)_d$ similarly satisfies:

$$\gamma(\Delta^n)_d^2 \leq \max_{0\leq k\leq d} \mathcal{G}_{2k}^{(\frac{n-2}{2})}(1) = \max_{0\leq k\leq d} (1 + \frac{4k}{n-1}) \cdot \binom{2k+n-2}{2k}.$$

□

8.3.2. Selecting a univariate square. The final ingredient we need for the proof of our main theorems is the following result of Fang and Fawzi [**FF21**].

LEMMA 8.14 ([**FF21**], Theorem 6). *Let $n, d \in \mathbb{N}$. Then for every $r \geq 2(n+1)d$ there exists a univariate sum of squares $q(x) = \sum_{k=0}^{2r} \lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x)$ of degree $2r$ with $\lambda_0 = 1$ and:*

$$\sum_{k=1}^{d} |1 - 1/\lambda_k| \leq \frac{2(n+1)^2 d^2}{r^2}.$$

We outline how to obtain this result here, following the strategy of [**FF21**]. It is quite similar to what happens in Section 7.2. We emphasize again a connection to the upper bounds in a univariate setting which was only implicitely present in [**FF21**]. We also state the intermediary result Lemma 8.16, which we need to prove Theorem 8.3 and Theorem 8.4 in Section 8.4. The first step of the argument is to linearize the quantity $\sum_{k=1}^{d} |1 - 1/\lambda_k|$.

LEMMA 8.15. *Let $n, d, r \in \mathbb{N}$ and let $q(x) = \sum_{k=0}^{2r} \lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x)$ be a sum of squares. Assuming that $\lambda_0 = 1$ and $\lambda_k \geq 1/2$ for $1 \leq k \leq d$ we have:*

$$\sum_{k=1}^{d} |1 - 1/\lambda_k| \leq 2 \sum_{k=1}^{d} (1 - \lambda_k)$$

Proof. As q is nonnegative on $[-1,1]$, we know that $\lambda_k \leq \lambda_0 = 1$ by (8.10). As $\lambda_k \geq 1/2$ for each k, we have:

$$\sum_{k=1}^{d} |1 - 1/\lambda_k| = \sum_{k=1}^{d} \frac{|1-\lambda_k|}{\lambda_k} \geq 2\sum_{k=1}^{d} |1-\lambda_k| = 2\sum_{k=1}^{d}(1-\lambda_k).$$

□

It remains to choose a sum of squares q minimizing $\sum_{k=1}^{d}(1-\lambda_k)$. This turns out to reduce to analyzing a univariate instance of the upper bounds (2.6).

Lemma 8.16. *Let $n, d \in \mathbb{N}$. Then for every $r \geq d$ there exists a univariate sum of squares $q(x) = \sum_{k=0}^{2r} \lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x)$ of degree $2r$ with $\lambda_0 = 1$ and:*

$$\sum_{k=1}^{d}(1-\lambda_k) \leq \frac{(n+1)^2 d^2}{r^2}.$$

Proof. For a univariate sum of squares $q(x) = \sum_{k=0}^{2r} \lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x)$, the coefficients λ_k are equal to (see relation (8.9)):

$$\lambda_k = \int_{-1}^{1} \overline{\mathcal{G}}_k^{(\frac{n-2}{2})}(x) q(x) w_n(x) dx \quad (0 \leq k \leq 2r),$$

where $\overline{\mathcal{G}}_k^{(\frac{n-2}{2})}$ is the normalization of the Gegenbauer polynomial of degree k satisfying $\max_{-1\leq x\leq 1} |\overline{\mathcal{G}}_k^{(\frac{n-2}{2})}(x)| = \overline{\mathcal{G}}_k^{(\frac{n-2}{2})}(1) = 1$. Let $h(x) := d - \sum_{k=1}^{d} \overline{\mathcal{G}}_k^{(\frac{n-2}{2})}(x)$. Note that:

$$\int_{-1}^{1} h(x) q(x) w_n(x) dx = d - \sum_{k=1}^{d} \lambda_k.$$

Selecting q optimally is thus equivalent to solving the optimization problem:

$$\inf_{q \in \Sigma[x]_{2r}} \left\{ \int_{-1}^{1} h(x) q(x) w_n(x) dx : \int_{-1}^{1} q(x) w_n(x) dx = 1 \right\}. \tag{8.27}$$

We recognize the program (8.27) as the upper bound $\mathrm{ub}(h)_r$ of (2.6) on the minimum $h_{\min} = h(1) = 0$ on $[-1,1]$ w.r.t. the measure $d\mu(x) = w_n(x)dx$.

As we have mentioned several times, the behaviour of the upper bounds in this univariate setting is well-understood. It is known that for the *linear* polynomial $u(x) = x$, the bounds have error $u^{(r)} - u_{\min} = (\xi_{r+1} + 1)$, where ξ_{r+1} is the smallest root of the Gegenbauer polynomial $\mathcal{G}_{r+1}^{(\frac{n-2}{2})}$ of degree $r+1$ [**dKL20b**]. These roots satisfy $(\xi_{r+1}+1) = O(1/r^2)$. The error $h^{(r)} - h_{\min}$ may then be bounded by considering the linear Taylor estimate of h at 1, which satisfies $h(1) + h'(1)x \geq h(x)$ for all $x \in [-1,1]$. See [**FF21**] for details. □

8.3.3. Proof of Theorem 8.1 and Theorem 8.2. We have now gathered all tools required to prove our main results. First, let $\mathbf{X} = B^n$ and let f be a polynomial on $\mathbf{X}$ of degree d. Recall that we may assume w.l.o.g. that $f_{\min} = 0$ and that $\|f\|_{\mathbf{X}} = 1$. We show how to construct an operator $\mathbf{K}$ satisfying the properties (P1), (P2) and (P3) for appropriate $\varepsilon > 0$, whose existence will immediately imply Theorem 8.1 by Lemma 8.5. See also Figure 8.1.

For $r \geq 2(n+1)d$, we select a univariate sum of squares:

$$q(x) = \sum_{k=0}^{2r} \lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x)$$

as in Lemma 8.14, i.e., such that $\lambda_0 = 1$ and $\sum_{k=1}^{d} |1 - 1/\lambda_k|$ is small. Consider the kernel $\mathrm{K}(\mathbf{x}, \mathbf{y}) := \mathrm{CD}_{2r}(\mathbf{x}, \mathbf{y}; \lambda)$ of (8.20). By Lemma 8.9, we know that the operator $\mathbf{K}$ associated to K satisfies (P2). Furthermore, Lemma 8.12 tell us that $\mathbf{K}$ satisfies (P1) and that it satisfies (P3) with:

$$\varepsilon = \max_{1 \leq k \leq d} \|f_k\|_{\mathbf{X}} \cdot \sum_{k=1}^{d} |1 - 1/\lambda_k| \leq \gamma(B^n)_d \cdot \frac{2(n+1)^2 d^2}{r^2}.$$

Here, we use (8.25) and Lemma 8.14 for the inequality. We may thus apply Lemma 8.5 to conclude the statement of Theorem 8.1 with constant:

$$C_B(n,d) = 2(n+1)^2 d^2 \gamma(B^n)_d. \tag{8.28}$$

In light of Proposition 8.13, this constant has the promised polynomial dependence on n (for fixed d) and on d (for fixed n).

The proof of Theorem 8.2 for $\mathbf{X} = \Delta^n$ is nearly identical. The only difference is that we should now select a sum of squares $q(x) = \sum_{k=0}^{4r} \lambda_k \mathcal{G}_k^{(\frac{n-2}{2})}(x)$ of degree $4r$ by applying Lemma 8.14 for $d \leftarrow 2d, r \leftarrow 2r$ and consider the kernel K defined in (8.22). The associated operator satisfies (P2) by Lemma 8.10. By Lemma 8.12, it satisfies (P1) and (P3) with:

$$\varepsilon \leq \gamma(\Delta^n)_d \cdot \frac{2(n+1)^2(2d)^2}{(2r)^2} = \gamma(\Delta^n)_d \cdot \frac{2(n+1)^2 d^2}{r^2},$$

using (8.24), Lemma 8.14 and (8.21). We may apply Lemma 8.5 to conclude the statement of Theorem 8.2 with:

$$C_\Delta(n,d) = 2(n+1)^2 d^2 \gamma(\Delta^n)_d, \tag{8.29}$$

which also has the required dependencies on n and d by Proposition 8.13.

8.4. The upper bounds

We apply the method of Section 6.2 to extend our results to the hierarchies of upper bounds $\mathrm{lb}(f, \mathcal{T}(\mathbf{X}))_r$ on the unit ball and the simplex. Recall that in

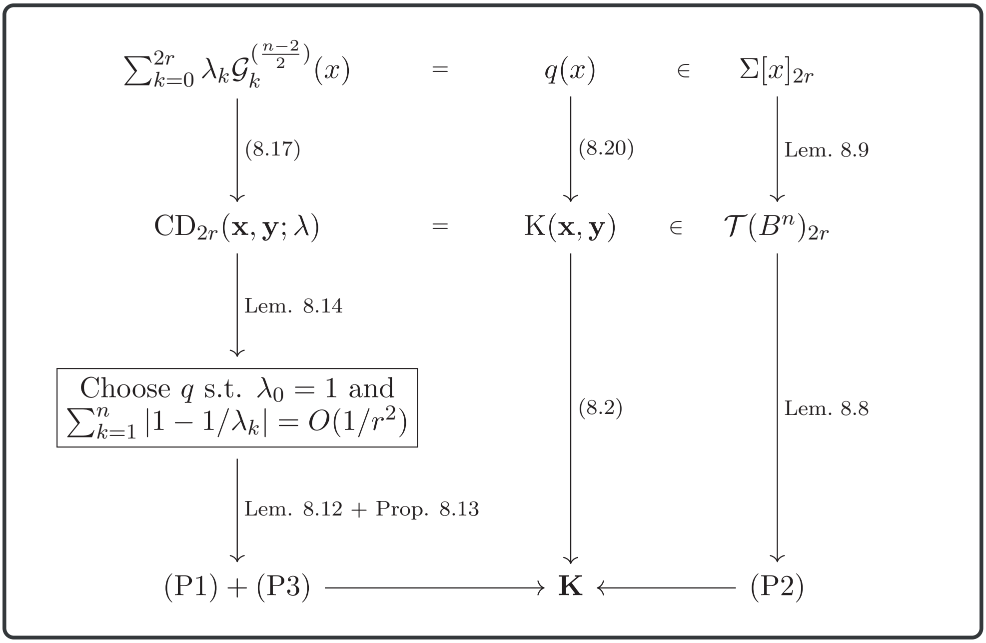

FIGURE 8.1. Overview of the construction of a a linear operator $\mathbf{K} : \mathbb{R}[\mathbf{x}]_d \to \mathbb{R}[\mathbf{x}]_d$ satisfying the properties (P1), (P2), (P3) of Section 8 for the unit ball. The construction for the standard simplex is analogous.

order to analyze these bounds, it is enough to exhibit an explicit probability density $\sigma \in \mathbb{R}[\mathbf{x}]$ on $(\mathbf{X}, \mu)$ which lies in $\mathcal{T}(\mathbf{X})_{2r}$, and for which the difference:

$$\int_{\mathbf{X}} f(\mathbf{x})\sigma(\mathbf{x})d\mu(\mathbf{x}) - f_{\min}$$

can be bounded from above. We exhibit such a σ here for $\mathbf{X} = B^n$, $\mu = \mu_B$ based on the perturbed Christoffel-Darboux kernels constructed in Section 8.2.

Write $\mathrm{K}(\mathbf{x},\mathbf{y}) = \mathrm{CD}_{2r}(\mathbf{x},\mathbf{y};\lambda)$ for such a kernel, where λ is chosen as in Lemma 8.16. Let $\mathbf{x}^* \in \mathbf{X}$ be a global minimizer of f over $\mathbf{X}$, and consider the polynomial σ given by:

$$\sigma(\mathbf{x}) = \mathrm{K}(\mathbf{x},\mathbf{x}^*) \quad (\mathbf{x} \in \mathbf{X}).$$

By Lemma 8.9, $\sigma \in \mathcal{T}(\mathbf{X})_{2r}$. For any polynomial $p \in \mathbb{R}[\mathbf{x}]_d$, we have:

$$\int_{\mathbf{X}} \sigma(\mathbf{x})p(\mathbf{x})d\mu(\mathbf{x}) = \mathbf{K}p(\mathbf{x}^*) = \sum_{k=0}^{d} \lambda_k p_k(\mathbf{x}^*) \quad (p_k \in H_k),$$

by definition. Therefore, as $\lambda_0 = 1$, σ is a probability density on $\mathbf{X}$. The polynomial σ is thus a feasible solution to the program (2.8) defining the bound $\mathrm{lb}(f, \mathcal{T}(\mathbf{X}))_r$. Furthermore, we have:

$$\int_{\mathbf{X}} f(\mathbf{x})\sigma(\mathbf{x})d\mu(\mathbf{x}) - f_{\min} = \mathbf{K}f(\mathbf{x}^*) - f(\mathbf{x}^*) \leq \sum_{k=1}^{d} |(1-\lambda_k) f_k(\mathbf{x}^*)|$$
$$\leq \gamma(\mathbf{X})_d \cdot \sum_{k=1}^{d} |1-\lambda_k| \leq \gamma(\mathbf{X})_d \cdot \frac{(n+1)^2 d^2}{r^2}.$$

We find that $\mathrm{lb}(f, \mathcal{T}(\mathbf{X}))_r - f_{\min} \leq \gamma(\mathbf{X})_d \cdot (n+1)^2 d^2 / r^2$. Theorem 8.3 follows immediately. For $\mathbf{X} = \Delta^n$, an analogous construction yields Theorem 8.4.

8.5. Discussion

We have shown a convergence rate in $O(1/r^2)$ for the Schmüdgen-type hierarchy of lower bounds $\mathrm{lb}(f, \mathcal{T}(\mathbf{X}))_r$ for the minimization of a polynomial f over the unit ball or simplex. Our result matches the best known rates for the hypersphere [**FF21**] and the hypercube (see Chapter 9). As a side result, we show similar convergence rates for the upper bounds (8.1) on these sets as well (w.r.t. the measures μ_B and μ_Δ). We repeat that convergence rates in $O(1/r^2)$ for the upper bounds on B^n and Δ^n were already available (but w.r.t. different reference measures).

Putinar- vs. Schmüdgen-type certificates. In light of the recent result [**BM21**], there is no longer a (large) theoretical gap between the best known convergence rates for the Putinar- and Schmüdgen-type hierarchies (see also Chapter 2). On the other hand, specialized convergence result for the Putinar-type bound $\mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r$ are so far available only in those cases where $\mathcal{Q}(\mathbf{X})_{2r} = \mathcal{T}(\mathbf{X})_{2r}$, i.e. where the Putinar- and Schmüdgen-type bounds coincide (which is the case on the binary hypercube, the hypersphere and the unit ball). It is an interesting open question whether specialized results for $\mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r$ can be shown in non-trivial cases as well, for instance on the simplex and hypercube (see Chapter 9). It seems unclear whether the techniques of the present chapter may be applied to the Putinar-type bounds on the simplex. This would require an analog of Lemma 8.10 showing membership of the kernel in the *quadratic module* rather than the preordering. Such an analog seems difficult to prove using the representation (8.22), and no obvious other representation is available.

The harmonic constant. We have shown in Proposition 8.13 that $\gamma(B^n)_d$ and $\gamma(\Delta^n)_d$ depend polynomially on n (for fixed d) and on d (for fixed n). As we have seen, the corresponding constants $\gamma(S^{n-1})_d$ and $\gamma(\{0,1\}^n)_d$ for the sphere and binary cube may actually be bounded *independently* of the dimension n for fixed $d \in \mathbb{N}$. It is an interesting open question whether this is true for $\gamma(B^n)_d$ and $\gamma(\Delta^n)_d$ as well. In a forthcoming paper, we will study the constants $\gamma(\mathbf{X})_d$ in more generality, focusing in particular on its asymptotic properties.

Acknowledgments. We wish to thank Monique Laurent for helpful suggestions on the presentation of the material, and Fernando Oliviera for fruitful discussions about kernels on the unit ball.

CHAPTER 9

Application: The unit box $[-1, 1]^n$

Intelligence is the ability of a living creature to perform pointless or unnatural acts.

Arkady Strugatsky, Roadside Picnic

This chapter is based on my joint work [**LS21**] *with Monique Laurent.*

Consider the problem of computing the global minimum:

$$f_{\min} := \min_{\mathbf{x} \in [-1,1]^n} f(\mathbf{x})$$

of a polynomial f of degree $d \in \mathbb{N}$ over the unit box $[-1, 1]^n \subseteq \mathbb{R}^n$. The set $[-1, 1]^n$ has the semialgebraic description:

$$[-1, 1]^n = \{\mathbf{x} \in \mathbb{R}^n : g_i(\mathbf{x}) := (1 - \mathbf{x}_i^2) \geq 0 \quad \forall i \in [n]\}.$$

The (truncated) *preordering* $\mathcal{T}([-1, 1]^n)_r$ is then defined as:

$$\mathcal{T}([-1, 1]^n)_r := \{\sum_{J \subseteq [n]} \sigma_J g_J : \sigma_J \in \Sigma[\mathbf{x}],\ \deg(\sigma_J g_J) \leq r\} \quad (g_J := \prod_{j \in J} g_j).$$

The preordering satisfies $\mathcal{T}([-1, 1]^n)_r \subseteq \mathcal{P}_+([-1, 1]^n)$ for all $r \in \mathbb{N}$, where $\mathcal{P}_+([-1, 1]^n)$ denotes the cone of polynomials nonnegative on $[-1, 1]^n$. For ease of exposition, we depart slightly from our earlier notation in this chapter and define the Schmüdgen-type lower bounds on $f_{\min}$ as:

$$\mathrm{lb}(f, \mathcal{T}([-1, 1]^n))_r := \sup\{\lambda \in \mathbb{R} : f - \lambda \in \mathcal{T}([-1, 1]^n)_r\} \tag{9.1}$$

(note that usually, we allow $f - \lambda \in \mathcal{T}([-1, 1]^n)_{2r}$). The reason for this small change is our later application of Theorem 9.8 below. As this is the only type of lower bound we consider in this chapter, we shall also write simply $\mathrm{lb}(f)_r = \mathrm{lb}(f, \mathcal{T}([-1, 1]^n))_r$. By definition, we have $f_{\min} \geq \mathrm{lb}(f)_{r+1} \geq \mathrm{lb}(f)_r$ for all $r \in \mathbb{N}$. Furthermore, we have $\lim_{r \to \infty} \mathrm{lb}(f)_r = f_{\min}$, which follows directly from the following special case of *Schmüdgen's Positivstellensatz.*

THEOREM 9.1 (Special case of Schmüdgen's Positivstellensatz [**Sch91**]). *Let $f \in \mathcal{P}_+([-1, 1]^n)$ be a polynomial. Then for any $\eta > 0$ there exists an $r \in \mathbb{N}$ such that $f + \eta \in \mathcal{T}([-1, 1]^n)_r$.*

Outline. In this chapter, we show a bound on the convergence rate of the lower bounds $\mathrm{lb}(f)_r$ to the global minimum $f_{\min}$ of f over $[-1,1]^n$ in $O(1/r^2)$. Alternatively, our result can be interpreted as a bound on the degree r in Schmüdgen's Positivstellensatz of the order $O(1/\sqrt{\eta})$ of a positivity certificate for $f+\eta$ when $f \in \mathcal{P}_+([-1,1]^n)$.

THEOREM 9.2. *Let f be a polynomial of degree $d \in \mathbb{N}$. Then there exists a constant $C(n,d) > 0$, depending only on n and d, such that:*

$$f_{\min} - \mathrm{lb}(f)_{(r+1)n} \leq \frac{C(n,d)}{r^2} \cdot (f_{\max} - f_{\min}) \quad \text{for all } r \geq \pi d\sqrt{2n}.$$

Furthermore, the constant $C(n,d)$ may be chosen such that it either depends polynomially on n (for fixed d) or it depends polynomially on d (for fixed n), see relation (9.13) for details.

COROLLARY 9.3. *Let $f \in \mathcal{P}_+([-1,1]^n)$ with degree d. Then, for any $\eta > 0$, we have:*

$$f+\eta \in \mathcal{T}([-1,1]^n)_{(r+1)n} \text{ for all } r \geq \max\left\{\pi d\sqrt{2n}, \frac{1}{\sqrt{\eta}}\sqrt{C(n,d)(f_{\max}-f_{\min})}\right\},$$

where $C(n,d)$ is the constant from Theorem 9.2. Hence we have $f+\eta \in \mathcal{T}([-1,1]^n)_r$ for $r = O(1/\sqrt{\eta})$.

PROOF. Let $\eta > 0$ and set $C_f := C(n,d) \cdot (f_{\max} - f_{\min})$. Pick an integer $r \geq \max\{\pi d\sqrt{2n}, \sqrt{C_f/\eta}\}$. Then we have:

$$f+\eta = \underbrace{f - \mathrm{lb}(f)_{(r+1)n}}_{\in \mathcal{T}([-1,1]^n)_{(r+1)n}} + \Big(\underbrace{\mathrm{lb}(f)_{(r+1)n} - f_{\min} + \frac{C_f}{r^2}}_{\geq 0 \text{ by Theorem 9.2}}\Big) + \underbrace{f_{\min}}_{\geq 0} + \Big(\underbrace{\eta - \frac{C_f}{r^2}}_{\geq 0}\Big),$$

which shows $f+\eta \in \mathcal{T}([-1,1]^n)_{(r+1)n}$. □

Overview of the proof. Let $f \in \mathbb{R}[\mathbf{x}]$ be a polynomial of degree d. To simplify our arguments and notation, we will work with the scaled function:

$$F := \frac{f - f_{\min}}{f_{\max} - f_{\min}},$$

for which $F_{\min} = 0$ and $F_{\max} = 1$. Since the inequality (9.2) is invariant under a positive scaling of f and adding a constant, it indeed suffices to show the result for the function F. To prove Theorem 9.2, we apply the method of Chapter 6. We outline the specialization of the method to this setting for convenience.

Let $\varepsilon > 0$ and consider the polynomial $\tilde{F} := F + \varepsilon$. Let $r \geq d$. Suppose that we are able to construct a (nonsingular) linear operator $\mathbf{K}_r : \mathbb{R}[\mathbf{x}]_r \to \mathbb{R}[\mathbf{x}]_r$

which has the following two properties:

$$\mathbf{K}_r(1) = 1, \tag{P1}$$

$$\mathbf{K}_r p \in \mathcal{T}([-1,1]^n)_{(r+1)n} \quad \text{for all } p \in \mathcal{P}_+([-1,1]^n)_d, \tag{P2}$$

$$\|\mathbf{K}_r^{-1}\tilde{F} - \tilde{F}\|_\infty := \max_{\mathbf{x}\in[-1,1]^n} |\mathbf{K}_r^{-1}\tilde{F}(\mathbf{x}) - \tilde{F}(\mathbf{x})| \leq \varepsilon. \tag{P3}$$

Then, by (P3), we have $\mathbf{K}_r^{-1}\tilde{F} \in \mathcal{P}_+([-1,1]^n)_d$. Indeed, as F is nonnegative on $[-1,1]^n$, $\tilde{F}(\mathbf{x}) = F(\mathbf{x}) + \varepsilon$ is greater than or equal to ε for all $\mathbf{x} \in [-1,1]^n$, and so (P3) tells us that after application of the operator $\mathbf{K}_r^{-1}$, the resulting polynomial $\mathbf{K}_r^{-1}\tilde{F}$ is nonnegative on $[-1,1]^n$. Using (P2), we may then conclude that $\tilde{F} = \mathbf{K}_r(\mathbf{K}_r^{-1}\tilde{F}) \in \mathcal{T}([-1,1]^n)_{(r+1)n}$. It follows that $F_{\min} - F_{((r+1)n)} \leq \varepsilon$ and thus $f_{\min} - \mathrm{lb}(f)_{(r+1)n} \leq \varepsilon \cdot (f_{\max} - f_{\min})$. We collect this in the next lemma for future reference.

LEMMA 9.4. *Assume that for some $r \geq d$ and $\varepsilon > 0$ there exists a non-singular operator $\mathbf{K}_r : \mathbb{R}[\mathbf{x}]_r \to \mathbb{R}[\mathbf{x}]_r$ which satisfies the properties* (P2) *and* (P3)*. Then we have:*

$$f_{\min} - \mathrm{lb}(f)_{(r+1)n} \leq \varepsilon \cdot (f_{\max} - f_{\min}).$$

In what follows, we will construct such an operator $\mathbf{K}_r$ for each $r \geq \pi d\sqrt{2n}$ and the parameter $\varepsilon := C(n,d)/r^2$, where the constant $C(n,d)$ will be specified later. Our main Theorem 9.2 then follows after applying Lemma 9.4.

We make use of the *polynomial kernel method* for our construction: after choosing a suitable kernel $\mathrm{K}_r : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$, we define the linear operator $\mathbf{K}_r : \mathbb{R}[\mathbf{x}]_r \to \mathbb{R}[\mathbf{x}]_r$ via the integral transform:

$$\mathbf{K}_r p(\mathbf{x}) := \int_{[-1,1]^n} \mathrm{K}_r(\mathbf{x},\mathbf{y}) p(\mathbf{y}) d\mu(\mathbf{y}) \quad (p \in \mathbb{R}[\mathbf{x}]_r).$$

Here, μ is the *Chebyshev measure* on $[-1,1]^n$ as defined in (9.2) below. A good choice for the kernel K_r is a multivariate version (see Section 9.2) of the well-known *Jackson kernel* $\mathrm{K}_r^{\mathrm{jac}}$ of degree r (see Section 9.1.3). For this choice of kernel, the operator $\mathbf{K}_r$ naturally satisfies (P2) (see Section 9.2.1). Furthermore, it diagonalizes with respect to the basis of $\mathbb{R}[\mathbf{x}]$ given by the (multivariate) *Chebyshev polynomials* (see Section 9.1.2). Property (P3) can then be verified by analyzing the eigenvalues of $\mathbf{K}_r$, which are closely related to the expansion of $\mathrm{K}_r^{\mathrm{jac}}$ in the basis of (univariate) Chebyshev polynomials (see Section 9.3). We end this section by illustrating our method of proof with a small example.

EXAMPLE 9.5. *Consider the polynomial $f(x) = 1 - x^2 - x^3 + x^4$, which is nonnegative on $[-1,1]$. For $r \in \mathbb{N}$, let $\mathbf{K}_r$ be the operator associated to the univariate Jackson kernel* (9.6) *of degree r, which satisfies* (P2) *(see Section 9.2.1). For $\eta = 0.1$, we observe that applying $\mathbf{K}_7^{-1}$ to $f + \eta$ yields a nonnegative function on $[-1,1]$, whereas applying $\mathbf{K}_5^{-1}$ does not (see Figure 9.1). Applying*

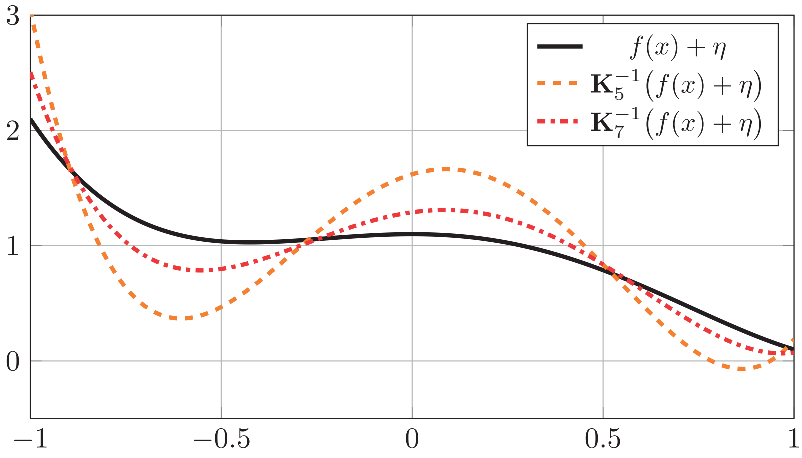

FIGURE 9.1. The polynomial $f(x)+\eta$ of Example 9.5 and its transformations under the inverse operators $\mathbf{K}_5^{-1}$ and $\mathbf{K}_7^{-1}$ associated to the Jackson kernels of degree 5 and 7.

the arguments of Section 9, we may thus conclude that $f+\eta \in \mathcal{T}([-1,1]^n)_8$, *but not that* $f+\eta \in \mathcal{T}([-1,1]^n)_6$.

9.1. Preliminaries

9.1.1. Notation. We write $\mathbb{R}[x]$ for the univariate polynomial ring, while reserving the bold-face notation $\mathbb{R}[\mathbf{x}] = \mathbb{R}[\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n]$ to denote the ring of polynomials in n variables. Similarly, $\Sigma[x] \subseteq \mathbb{R}[x]$ and $\Sigma[\mathbf{x}] \subseteq \mathbb{R}[\mathbf{x}]$ denote the sets of univariate and n-variate sum-of-squares polynomials, respectively, consisting of all polynomials of the form $p = p_1^2 + p_2^2 + \cdots + p_m^2$ for certain polynomials $p_1, \ldots, p_m$ and $m \in \mathbb{N}$. For a polynomial $p \in \mathbb{R}[\mathbf{x}]$, we write $p_{\min}, p_{\max}$ for its minimum and maximum over $[-1,1]^n$, respectively, and $\|p\|_\infty := \sup_{\mathbf{x}\in[-1,1]^n} |p(\mathbf{x})|$ for its sup-norm on $[-1,1]^n$.

9.1.2. Chebyshev polynomials. We recall some properties of the Chebyshev polynomials, introduced in Chapter 1. Let μ be the normalized *Chebyshev measure* on $[-1,1]^n$, defined by:

$$d\mu(\mathbf{x}) = \frac{d\mathbf{x}_1}{\pi\sqrt{1-\mathbf{x}_1^2}} \cdots \frac{d\mathbf{x}_n}{\pi\sqrt{1-\mathbf{x}_n^2}}. \tag{9.2}$$

Note that μ is a probability measure on $[-1,1]^n$, meaning that $\int_{[-1,1]^n} d\mu = 1$. We write $\langle \cdot, \cdot \rangle_\mu$ for the corresponding inner product on $\mathbb{R}[\mathbf{x}]$, given by:

$$\langle f, g \rangle_\mu := \int_{[-1,1]^n} f(\mathbf{x}) g(\mathbf{x}) d\mu(\mathbf{x}).$$

For $k \in \mathbb{N}$, let $\mathcal{C}_k$ be the univariate *Chebyshev polynomial* of degree k, defined by:

$$\mathcal{C}_k(\cos\theta) := \cos(k\theta) \quad (\theta \in \mathbb{R}).$$

Note that $|\mathcal{C}_k(x)| \leq 1$ for all $x \in [-1,1]$ and that $\mathcal{C}_0 = 1$. The Chebyshev polynomials satisfy the orthogonality relations:

$$\langle \mathcal{C}_a, \mathcal{C}_b \rangle_\mu = \int_{-1}^{1} \mathcal{C}_a(x)\mathcal{C}_b(x) d\mu(x) = \begin{cases} 0 & a \neq b, \\ 1 & a = b = 0, \\ \frac{1}{2} & a = b \neq 0. \end{cases} \tag{9.3}$$

A univariate polynomial p may therefore be expanded as:

$$p = p_0 + \sum_{k=1}^{\deg(p)} 2p_k\mathcal{C}_k, \quad \text{where } p_k := \langle \mathcal{C}_k, p \rangle_\mu.$$

For $\kappa \in \mathbb{N}^n$, we consider the *multivariate* Chebyshev polynomial $\mathcal{C}_\kappa$, defined by setting:

$$\mathcal{C}_\kappa(\mathbf{x}) := \prod_{i=1}^{n} \mathcal{C}_{\kappa_i}(\mathbf{x}_i).$$

The multivariate Chebyshev polynomials form a basis for $\mathbb{R}[\mathbf{x}]$ and satisfy the orthogonality relations:

$$\langle \mathcal{C}_\alpha, \mathcal{C}_\beta \rangle_\mu = \int_{[-1,1]^n} \mathcal{C}_\alpha(\mathbf{x})\mathcal{C}_\beta(\mathbf{x}) d\mu(\mathbf{x}) = \begin{cases} 0 & \alpha \neq \beta, \\ 1 & \alpha = \beta = 0, \\ 2^{-w(\alpha)} & \alpha = \beta \neq 0. \end{cases} \tag{9.4}$$

Here, $w(\alpha) := |\{i \in [n] : \alpha_i \neq 0\}|$ denotes the Hamming weight of $\alpha \in \mathbb{N}^n$.

We use the notation $\mathbb{N}^n_d \subseteq \mathbb{N}^n$ to denote the set of n-tuples $\alpha \in \mathbb{N}^n$ with $|\alpha| = \sum_{i=1}^n \alpha_i \leq d$. As in the univariate case, we may expand any n-variate polynomial p as:

$$p = \sum_{\kappa \in \mathbb{N}^n_{\deg(p)}} 2^{w(\kappa)} p_\kappa \mathcal{C}_\kappa, \quad \text{where } p_\kappa := \langle \mathcal{C}_\kappa, p \rangle_\mu. \tag{9.5}$$

9.1.3. The Jackson kernel. For $r \in \mathbb{N}$ and for coefficients $\lambda_k^{(r)} \in \mathbb{R}$ to be specified below in (9.7), consider the kernel $\mathrm{K}_r^{\mathrm{jac}} : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ given by:

$$\mathrm{K}_r^{\mathrm{jac}}(x,y) := 1 + 2\sum_{k=1}^{r} \lambda_k^{(r)} \mathcal{C}_k(x)\mathcal{C}_k(y). \tag{9.6}$$

Note that $\mathrm{K}_r^{\mathrm{jac}}(x,y) = \mathrm{CD}(x,y;\lambda^{(r)})$ thus equals the perturbed Christoffel-Darboux kernel for $[-1,1]$ w.r.t. the Chebyshev measure defined in (6.7). We associate a linear operator $\mathbf{K}_r^{\mathrm{jac}} : \mathbb{R}[x]_r \to \mathbb{R}[x]_r$ to this kernel by setting:

$$\mathbf{K}_r^{\mathrm{jac}} p(x) := \int_{-1}^{1} \mathrm{K}_r^{\mathrm{jac}}(x,y)p(y)d\mu(y) \quad (p \in \mathbb{R}[x]_r).$$

Using the orthogonality relations (9.3), and writing $\lambda_0^r := 1$, we see that:

$$\mathbf{K}_r^{\text{jac}}\mathcal{C}_k(x) := \int_{-1}^{1} \mathrm{K}_r^{\text{jac}}(x,y)\mathcal{C}_k(y)d\mu(y) = \lambda_k^r\mathcal{C}_k(x) \quad (0 \le k \le r).$$

In other words, $\mathbf{K}_r^{\text{jac}}$ is a diagonal operator with respect to the Chebyshev basis of $\mathbb{R}[x]_r$, and its eigenvalues are given by $\lambda_0^r = 1, \lambda_1^{(r)}, \ldots, \lambda_r^{(r)}$. In what follows, we set:

$$\lambda_k^{(r)} = \frac{1}{r+2}\big((r+2-k)\cos(k\theta_r) + \frac{\sin(k\theta_r)}{\sin(\theta_r)}\cos(\theta_r)\big) \quad (1 \le k \le r), \tag{9.7}$$

with $\theta_r = \frac{\pi}{r+2}$. We then obtain the so-called *Jackson kernel* (see, e.g., [**WWAF06**]). The following properties of the Jackson kernel are crucial to our analysis.

PROPOSITION 9.6. *For every $d, r \in \mathbb{N}$ with $d \le r$, we have:*

(i) $\mathrm{K}_r^{\text{jac}}(x,y) \ge 0$ *for all* $x, y \in [-1,1]$,
(ii) $1 \ge \lambda_k^{(r)} > 0$ *for all* $0 \le k \le r$, *and*
(iii) $|1-\lambda_k^{(r)}| = 1-\lambda_k^{(r)} \le \frac{\pi^2 d^2}{(r+2)^2}$ *for all* $0 \le k \le d$.

PROOF. Nonnegativity of the Jackson kernel is a well-known fact, and is verified, e.g., in [**dKHL17**]. We check that the other properties (ii)-(iii) hold as well.

Second property (ii): Note that when $k \le (r+2)/2$, both terms of (9.7) are positive, and so certainly $\lambda_k^{(r)} > 0$. So assume $(r+2)/2 < k \le r$. Set $h = r+2-k$, so that $k\theta_r = \pi - h\theta_r$, $2 \le h < (r+2)/2$, and

$$(r+2)\lambda_k^{(r)} = -h\cos(h\theta_r) + \frac{\sin(h\theta_r)}{(\sin\theta_r)}\cos(\theta_r) \tag{9.8}$$

It remains to show that the RHS of (9.8) is positive for all $2 \le h < (r+2)/2$. Note that $1 > \cos(\theta_r) > 0$, $\sin(\theta_r) \ge 0$ and that $\sin(h\theta_r) \ge 0$ for all $2 \le h < (r+2)/2$. We proceed by induction. For $h = 2$, we compute:

$$\begin{aligned}-h\cos(h\theta_r) + \frac{\sin(h\theta_r)}{\sin(\theta_r)}\cos(\theta_r) &= -2(2\cos^2(\theta_r)-1) + 2\cos^2(\theta_r)\\ &= -2\cos^2(\theta_r) + 2 > 0,\end{aligned}$$

which settles the base of induction. For $h \geq 2$, we compute:

$$\begin{aligned}
&-(h+1)\cos((h+1)\theta_r) + \sin((h+1)\theta_r)\frac{\cos(\theta_r)}{\sin(\theta_r)} \\
&\quad = -(h+1)\big(\cos(h\theta_r)\cos(\theta_r) - \sin(h\theta_r)\sin(\theta_r)\big) \\
&\qquad + \big(\sin(h\theta_r)\cos(\theta_r) + \cos(h\theta_r)\sin(\theta_r)\big)\frac{\cos(\theta_r)}{\sin(\theta_r)} \\
&\quad = -h\cos(h\theta_r)\cos(\theta_r) + (h+1)\sin(h\theta_r)\sin(\theta_r) + \frac{\sin(h\theta_r)}{\sin(\theta_r)}\cos^2(\theta_r) \\
&\quad = \underbrace{\cos(\theta_r)}_{>0}\underbrace{\Big(-h\cos(h\theta_r) + \frac{\sin(h\theta_r)}{\sin(\theta_r)}\cos(\theta_r)\Big)}_{\geq 0 \text{ by the induction assumption}} + (h+1)\underbrace{\sin(h\theta_r)\sin(\theta_r)}_{\geq 0} \\
&\quad \geq 0.
\end{aligned}$$

We conclude that $\lambda_k^{(r)} > 0$ for all $k \in [r]$. To see that $\lambda_k^{(r)} \leq 1$, note that for all $k \in \mathbb{N}$, $\mathcal{C}_k(x) \leq 1$ for $-1 \leq x \leq 1$ and $\mathcal{C}_k(1) = 1$. We can thus compute:

$$\lambda_k^{(r)} = \lambda_k^{(r)}\mathcal{C}_k(1) = \int_{-1}^{1} \mathrm{K}_r^{\mathrm{jac}}(1,y)\mathcal{C}_k(y)d\mu(y) \leq \int_{-1}^{1} \mathrm{K}_r^{\mathrm{jac}}(1,y)d\mu(y) = \lambda_0^{(r)} = 1,$$

making use of the nonnegativity of $\mathrm{K}_r^{\mathrm{jac}}(x,y)$ on $[-1,1]^2$ for the inequality.

Third property (iii): Using the expression of λ_r^k in (9.7) we have

$$1 - \lambda_r^k = 1 - \frac{r+2-k}{r+2}\cos(k\theta_r) - \frac{1}{r+2}\frac{\sin(k\theta_r)\cos(\theta_r)}{\sin(\theta_r)}.$$

We now bound each trigonometric term using the fact that:

$$\cos(x) \geq 1 - \frac{1}{2}x^2, \quad x - \frac{1}{6}x^3 \leq \sin(x) \leq x \quad (x \in \mathbb{R}). \tag{9.9}$$

When $k = 1$ we immediately get:

$$1 - \lambda_r^1 = 1 - \cos(\theta_r) \leq \frac{1}{2}\theta_r^2 = \frac{\pi^2}{2(r+2)^2} \leq \frac{d^2\pi^2}{(r+2)^2}.$$

Assume now $2 \leq k \leq d$. Using (9.9) combined with $\cos(\theta_r), \sin(\theta_r), \sin(k\theta_r) > 0$ we obtain:

$$\frac{\sin(k\theta_r)\cos(\theta_r)}{\sin(\theta_r)} \geq \big(k\theta_r - \frac{1}{6}k^3\theta_r^3\big)\big(1 - \frac{1}{2}\theta_r^2\big)\frac{1}{\theta_r} \geq k - \frac{k}{2}\theta_r^2\big(1 + \frac{k^2}{3}\big)$$

and thus:

$$\begin{aligned}1-\lambda_r^k &\le 1-\frac{r+2-k}{r+2}\Big(1-\frac{k^2\theta_r^2}{2}\Big)-\frac{1}{r+2}\Big(k-\frac{k}{2}\theta_r^2\big(1+\frac{k^2}{3}\big)\Big)\\ &= \underbrace{\frac{r+2-k}{r+2}}_{\le 1}\frac{k^2\theta_r^2}{2}+\underbrace{\frac{k}{2(r+2)}}_{\le 1/2}\theta_r^2\underbrace{\big(1+\frac{k^2}{3}\big)}_{\le\frac{2}{3}k^2 \text{ if } k\ge 2}\\ &\le k^2\theta_r^2\le\frac{d^2\pi^2}{(r+2)^2}.\end{aligned}$$

This concludes the proof if $k \geq 2$. □

9.2. Construction of the linear operator

As noted before, in order to prove Theorem 9.2 it suffices to construct a linear operator $\mathbf{K}_r : \mathbb{R}[\mathbf{x}]_r \to \mathbb{R}[\mathbf{x}]_r$ that is nonsingular and satisfies (P1), (P2) and (P3). For this purpose we define the multivariate Jackson kernel $\mathrm{K}_r^{\mathrm{jac}} : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ by setting:

$$\mathrm{K}_r^{\mathrm{jac}}(\mathbf{x},\mathbf{y}) := \prod_{i=1}^n \mathrm{K}_r^{\mathrm{jac}}(x_i, y_i),$$

where $\mathrm{K}_r^{\mathrm{jac}}$ is the (univariate) Jackson kernel from (9.6). Now let $\mathbf{K}_r$ be the corresponding kernel operator defined by:

$$\mathbf{K}_r p(\mathbf{x}) = \int_{\mathbf{y}\in[-1,1]^n} \mathrm{K}_r^{\mathrm{jac}}(\mathbf{x},\mathbf{y})p(\mathbf{y})d\mu(\mathbf{y}) \quad (p \in \mathbb{R}[\mathbf{x}]_r).$$

The operator $\mathbf{K}_r$ is diagonal w.r.t. the (multivariate) Chebyshev basis, and its eigenvalues can be expressed in terms of the coefficients $\lambda_k^{(r)}$ of the Jackson kernel, as the following lemma shows.

Lemma 9.7. *The operator $\mathbf{K}_r$ is diagonal w.r.t. the Chebyshev basis for $\mathbb{R}[\mathbf{x}]_r$, and its eigenvalues are given by:*

$$\lambda_\kappa^{(r)} := \prod_{i=1}^n \lambda_{\kappa_i}^{(r)} \quad (\kappa \in \mathbb{N}_r^n).$$

Proof. For $\kappa \in \mathbb{N}_r^n$, we see that:

$$\begin{aligned}\mathbf{K}_r\mathcal{C}_\kappa(\mathbf{x}) &= \int_{\mathbf{y}\in[-1,1]^n} \mathrm{K}_r^{\mathrm{jac}}(\mathbf{x},\mathbf{y})\mathcal{C}_\kappa(\mathbf{y})d\mu(\mathbf{y})\\ &= \prod_{i=1}^n\Big(\int_{y_i\in[-1,1]} \mathrm{K}_r^{\mathrm{jac}}(\mathbf{x}_i,\mathbf{y}_i)\mathcal{C}_{\kappa_i}(\mathbf{y}_i)d\mu(\mathbf{y}_i)\Big)\\ &= \prod_{i=1}^n \lambda_{\kappa_i}^{(r)}\mathcal{C}_{\kappa_i}(\mathbf{x}_i) = \lambda_\kappa^{(r)}\mathcal{C}_\kappa(\mathbf{x}),\end{aligned}$$

as required. □

Note that $\mathrm{K}^{\mathrm{jac}}(\mathbf{x},\mathbf{y}) = \mathrm{CD}(\mathbf{x},\mathbf{y};\lambda^{(r)})$ is thus the perturbed Christoffel-Darboux kernel for $[-1.1]^n$ w.r.t. the Chebyshev measure defined in (6.7). It follows immediately from Proposition 9.6(ii) that $\mathbf{K}_r$ has only nonzero eigenvalues and thus is non-singular. As $\lambda_0^{(r)} = 1$, we also have $\mathbf{K}_r(1) = 1$, meaning $\mathbf{K}_r$ satisfies (P1). We show that it further satisfies (P2) and (P3).

9.2.1. Verification of property (P2). Consider the following strengthening of Schmüdgen's Positivstellensatz in the univariate case.

THEOREM 9.8 (Fekete, Markov-Lukácz (see [**PR00**])). *Let p be a univariate polynomial of degree r, and assume that $p \geq 0$ on the interval $[-1,1]$. Then p admits a representation of the form:*

$$p(x) = \sigma_0(x) + \sigma_1(x)(1-x^2), \tag{9.10}$$

where $\sigma_0, \sigma_1 \in \Sigma[x]$ and σ_0 and $\sigma_1 \cdot (1-x^2)$ are of degree at most $r+1$. In other words, in view of (9), we have $p \in \mathcal{T}([-1,1])_{r+1}$.

By Proposition 9.6(i), for any $y \in [-1,1]$, the polynomial $x \mapsto \mathrm{K}_r^{\mathrm{jac}}(x,y)$ is nonnegative on $[-1,1]$ and therefore admits a representation of the form (9.10). This implies directly that the multivariate polynomial $\mathbf{x} \mapsto \mathrm{K}_r^{\mathrm{jac}}(\mathbf{x},\mathbf{y}) = \prod_{i=1}^n \mathrm{K}_r^{\mathrm{jac}}(\mathbf{x}_i,\mathbf{y}_i)$ belongs to $\mathcal{T}([-1,1]^n)_{(r+1)n}$ for all $\mathbf{y} \in [-1,1]^n$.

LEMMA 9.9 (Specialization of Lemma 6.2). *The operator $\mathbf{K}_r$ satisfies property* (P2)*, that is, we have $\mathbf{K}_r p \in \mathcal{T}([-1,1]^n)_{(r+1)n}$ for all $p \in \mathcal{P}_+([-1,1]^n)$.*

PROOF. One way to see this is as follows. Let $p \in \mathcal{P}_+([-1,1]^n)$ and let $\{\mathbf{y}_i : i \in [N]\} \subseteq [-1,1]^n$ and $w_i > 0$ $(i \in [N])$ form a quadrature rule for integration of polynomials up to degree $r + \deg(p)$ over $[-1,1]^n$; that is, $\int_{[-1,1]^n} q(\mathbf{x})d\mu(\mathbf{x}) = \sum_{i=1}^N w_i q(\mathbf{y}_i)$ for any $q \in \mathbb{R}[\mathbf{x}]_{r+\deg(p)}$. Then, we have $\mathbf{K}_r p(\mathbf{x}) = \sum_{i=1}^N \mathrm{K}_r(\mathbf{x},\mathbf{y}_i)p(\mathbf{y}_i)w_i$ with $p(\mathbf{y}_i)w_i \geq 0$ for all i, which shows that $\mathbf{K}_r p \in \mathcal{T}([-1,1]^n)_{(r+1)n}$. □

9.3. Analysis of the linear operator

We may decompose the polynomial $\tilde{F} = F + \varepsilon$ into the multivariate Chebyshev basis (9.5):

$$\tilde{F} = \varepsilon + \sum_{\kappa \in \mathbb{N}_d^n} 2^{w(\kappa)} F_\kappa \mathcal{C}_\kappa, \quad \text{where } F_\kappa = \langle F, T_\kappa \rangle_\mu.$$

By Lemma 9.7, we then have:

$$\begin{aligned} \|\mathbf{K}_r^{-1}\tilde{F} - \tilde{F}\|_\infty &= \|\sum_{\kappa \in \mathbb{N}_d^n} (1/\lambda_\kappa^{(r)}) 2^{w(\kappa)} F_\kappa \mathcal{C}_\kappa - 2^{w(\kappa)} F_\kappa \mathcal{C}_\kappa\|_\infty \\ &\leq \sum_{\kappa \in \mathbb{N}_d^n} 2^{w(\kappa)} |F_\kappa| |1 - 1/\lambda_\kappa^{(r)}|, \end{aligned} \tag{9.11}$$

making use of the fact that $\lambda_0 = 1$ and $|\mathcal{C}_\kappa(\mathbf{x})| \leq 1$ for all $\mathbf{x} \in [-1,1]^n$. It remains to analyze the expression at the right-hand side of (9.11). First, we bound the size of $|F_\kappa|$ for $\kappa \in \mathbb{N}^n$.

Lemma 9.10. *We have* $|F_\kappa| = |\langle F, \mathcal{C}_\kappa \rangle_\mu| \leq 2^{-w(\kappa)/2}$ *for all* $\kappa \in \mathbb{N}^n$.

Proof. Since μ is a probability measure on $[-1,1]^n$, we have $\|F\|_\mu \leq \|F\|_\infty \leq 1$. Using the Cauchy-Schwarz inequality and (9.4), we then find:

$$\langle F, \mathcal{C}_\kappa \rangle_\mu \leq \|F_\kappa\|_\mu \|\mathcal{C}_\kappa\|_\mu \leq \|\mathcal{C}_\kappa\|_\mu = 2^{-w(\kappa)/2}.$$

□

To bound the parameter $|1-1/\lambda_\kappa^{(r)}|$, we first prove a bound on $|1-\lambda_\kappa^{(r)}|$, which we obtain by applying Bernoulli's inequality.

Lemma 9.11 (Bernoulli's inequality). *For any* $x \in [0,1]$ *and* $t \geq 1$, *we have:*

$$1 - (1-x)^t \leq tx. \tag{9.12}$$

Lemma 9.12. *For any* $\kappa \in \mathbb{N}^n_d$ *and* $r \geq \pi d$, *we have:*

$$|1 - \lambda_\kappa^{(r)}| \leq \frac{n\pi^2 d^2}{r^2}.$$

Proof. By Proposition 9.6, we know that $0 \leq \gamma_k := (1-\lambda_k^{(r)}) \leq \pi^2 d^2/r^2 \leq 1$ for $0 \leq k \leq d$. Writing $\gamma := \max_{0 \leq k \leq d} \gamma_k$, we compute:

$$1 - \lambda_\kappa^{(r)} = 1 - \prod_{i=1}^n \lambda_{\kappa_i}^{(r)} = 1 - \prod_{i=1}^n (1-\gamma_{\kappa_i}) \leq 1 - (1-\gamma)^n \leq n\gamma \leq \frac{n\pi^2 d^2}{r^2},$$

making use of (9.12) for the second to last inequality. □

Lemma 9.13. *Assuming that* $r \geq \pi d\sqrt{2n}$, *we have:*

$$|1 - 1/\lambda_\kappa^{(r)}| \leq \frac{2n\pi^2 d^2}{r^2}.$$

Proof. Under the assumption, and using the previous lemma, we have $|1 - \lambda_\kappa^{(r)}| \leq 1/2$, which implies that $\lambda_\kappa^{(r)} \geq 1/2$. We may then bound:

$$|1 - 1/\lambda_\kappa^{(r)}| = |\frac{1-\lambda_\kappa^{(r)}}{\lambda_\kappa^{(r)}}| \leq 2|1 - \lambda_\kappa^{(r)}| \leq \frac{2n\pi^2 d^2}{r^2}.$$

□

Putting things together and using (9.11), Lemma 9.10 and Lemma 9.12 we find that:

$$\begin{aligned}\|\mathbf{K}_r^{-1}\tilde{F} - \tilde{F}\|_\infty &\leq \sum_{\kappa \in \mathbb{N}^n_d} 2^{w(\kappa)} |F_\kappa| |1 - 1/\lambda_\kappa^{(r)}| \\ &\leq \sum_{\kappa \in \mathbb{N}^n_d} 2^{w(\kappa)/2} \cdot \frac{2n\pi^2 d^2}{r^2} \leq |\mathbb{N}^n_d| \cdot \max_{\kappa \in \mathbb{N}^n_d} 2^{w(\kappa)/2} \cdot \frac{2n\pi^2 d^2}{r^2}.\end{aligned}$$

Hence $\mathbf{K}_r$ satisfies (P3) with $\varepsilon = C(n,d)/r^2$, where:

$$C(n,d) := |\mathbb{N}^n_d| \cdot \max_{\kappa \in \mathbb{N}^n_d} 2^{w(\kappa)/2} \cdot 2n\pi^2 d^2.$$

In view of Lemma 9.4, we have thus proven Theorem 9.2. Finally, we can bound the constant $C(n,d)$ in two ways. On the one hand, we have:

$$|\mathbb{N}^n_d| = \binom{n+d}{n} = \prod_{i=1}^{n} \frac{d+i}{i} \leq (d+1)^n \text{ and } \max_{\kappa \in \mathbb{N}^n_d} w(\kappa) \leq n,$$

resulting in a polynomial dependence of $C(n,d)$ on d for fixed n. On the other hand, we have:

$$|\mathbb{N}^n_d| = \binom{n+d}{d} \leq (n+1)^d \text{ and } \max_{\kappa \in \mathbb{N}^n_d} w(\kappa) \leq d,$$

resulting in a polynomial dependence of $C(n,d)$ on n for fixed d. Namely, we have:

$$C(n,d) \leq 2\pi^2 d^2 n 2^{n/2}(d+1)^n \text{ and } C(n,d) \leq 2\pi^2 d^2 n 2^{d/2}(n+1)^d. \tag{9.13}$$

9.4. Discussion

We have shown that the error of the degree $r+1$ Lasserre-type bound (9.1) for the minimization of a polynomial over the hypercube $[-1,1]^n$ is of the order $O(1/r^2)$ when using a sum-of-squares decomposition in the truncated preordering. Alternatively, if f is a polynomial nonnegative on $[-1,1]^n$ and $\eta > 0$, our result may be interpreted as showing a bound in $O(1/\sqrt{\eta})$ on the degree of a Schmüdgen-type certificate of positivity for $f+\eta$. The dependence on the dimension n and the degree d of f in the constants of our result is both polynomial in n (for fixed d), and polynomial in d (for fixed n).

Earlier results. A convergence rate in $O(1/r)$ for the lower bounds $\mathrm{lb}(f)_r$ was already known in the literature [**dKL10**]. This rate holds in fact for a weaker hierarchy of bounds obtained by restricting in (9.1) to decompositions of the polynomial $f - \lambda$ involving factors σ_J that are nonnegative scalars (instead of sums of squares), also known as Handelman-type decompositions (thus replacing the preordering $\mathcal{T}([-1,1]^n)_r$ by its subset H_r of polynomials having a Handelman-type decomposition). The analysis in [**dKL10**] relies on employing the *Bernstein operator* $\mathbf{B}_r$, which has the property of mapping a polynomial nonnegative over the hypercube to a polynomial in the set $H_{rn} \subseteq \mathcal{T}([-1,1]^n)_{rn}$.

An application of our main result. As we discussed in Section 2.3, Baldi & Mourrain [**BM21**] have recently improved the previously best-known convergence rate of [**NS07**] for the Putinar-type bounds on a *general* Archimedean semialgebraic set $\mathbf{X}$. Roughly speaking, their method of proof relies on embedding $\mathbf{X}$ in a box $[-R,R]^n$ of large enough size $R > 0$, and then relating positivity certificates on $\mathbf{X}$ to those on $[-R,R]^n$. Our present result on $[-1,1]^n$ then allows them to conclude their analysis. Their argument relies on the fact the constant $C(n,d)$ in Theorem 9.2 may be chosen to depend polynomially on the degree d of f. Such a dependence was not shown in the earlier work [**dKL10**].

The harmonic constant. A question left open in this work is whether it is possible to show Theorem 9.2 with a constant $C(d)$ that only depends on the degree d of f, and not on the number of variables n (cf. (9.13)). This question is motivated by the fact that for the analysis of the analogous hierarchies for the unit sphere in [**FF21**] and for the boolean hypercube in Chapter 7 the existence of such a constant (depending only on d) was in fact shown.

Improving upon the Jackson kernel. We choose to use the (multivariate) Jackson kernel in this chapter mostly because its relevant properties are well-understood and it is thus easy to analyze. In principle, there could be other kernels that would yield better convergence guarantees. In particular, it would be nice to have kernels that lie in the quadratic module $\mathcal{Q}([-1,1]^n)_r$ (instead of the preordering), thereby allowing an analysis of the Putinar-type bounds. In their recent work [**KdK22**], the authors show how to construct non-negative kernels numerically whose associated operators converge to the identity quickly. Their results could potentially be used to analyze the lower bounds. However, one would probably first have to find analytical expressions for their kernels (after which it might be possible to show membership in the preordering or quadratic module).

The upper bounds. As we explain in Section 6.2, our method of proof in this chapter would in principle allow us to obtain convergence rates for the *upper* bounds $\mathrm{ub}(f, \mathcal{T}([-1,1]^n), \mu)_r$ as well. In fact, this is precisely what Hess, de Klerk and Laurent do in [**dKHL17**]. There, they implicitely use the polynomial kernel method, and the Jackson kernel in particular, to obtain a convergence rate in $O(1/r^2)$ for these bounds.

Acknowledgments. We thank Lorenzo Baldi and Bernard Mourrain for their helpful suggestions. We also thank Etienne de Klerk and Felix Kirschner for useful discussions.

Part 3

Independent sets in geometric hypergraphs

CHAPTER 10

A recursive theta number for geometric hypergraphs

It seems all paths lead to ϑ!

Michel Goemans

This chapter is based on my joint works [**CSdOFSV21, CSdOFSV22**] *with Fernando de Oliveira, Davi Silva and Frank Vallentin.*

Let $G = (V, E)$ be a (finite) graph. A subset $S \subseteq V$ is called a *stable* (or *independent*) set if no two vertices $v, w \in S$ are adjacent, i.e., if $\{v, w\} \notin E$ for all $v, w \in S$. The *stability number* (or *independence number*) $\alpha(G)$ of G is equal to the largest cardinality of a stable set in G. Computing the stability number of a graph is a classical NP-hard problem [**Kar72**].

The celebrated *Lovász theta number* $\vartheta(G)$ of G satisfies $\alpha(G) \leq \vartheta(G) \leq \chi(\overline{G})$, where $\alpha(G)$ is the independence number of G and $\chi(\overline{G})$ is the chromatic number of the complement of G. That is, the edges of $\overline{G}$ are the non-edges of G and vice versa. Crucially, the theta number can be computed efficiently using semidefinite programming.

Originally, Lovász [**Lov79**] introduced ϑ to determine the Shannon capacity of the 5-cycle. The theta number turned out to be a versatile tool in optimization, with applications in combinatorics and geometry. It is related to spectral bounds like Hoffman's bound, as noted by Lovász in his paper (cf. Bachoc, DeCorte, Oliveira, and Vallentin [**BDdOFV14**]), and also to Delsarte's linear programming bound in coding theory, as observed independently by McEliece, Rodemich, and Rumsey [**MRR78**] and Schrijver [**Sch79**].

One way to define $\vartheta(G)$ is to consider the following formulation of $\alpha(G)$ as an integer program:

$$\alpha(G) = \max_{\mathbf{x} \in \{0,1\}^V} \{\sum_{i \in V} \mathbf{x}_i : \mathbf{x}_i \mathbf{x}_j = 0 \quad \forall \{i, j\} \in E\}. \tag{10.1}$$

The feasible solutions of (10.1) are precisely the *incidence vectors* $\mathbf{x}_S$ of stable sets $S \subseteq V$, defined by $(\mathbf{x}_S)_i = 1$ if $i \in S$ and 0 otherwise. For such a feasible solution $\mathbf{x}_S$, the positive semidefinite matrix $X = \mathbf{x}_S \mathbf{x}_S^\top / |S|$ satisfies $X_{ij} = 0$ for $\{i, j\} \in E$ and $\mathrm{Tr}(X) = 1$. Furthermore, if J is the all-ones matrix, we have $\langle X, J \rangle := \mathrm{Tr}(XJ) = \sum_{i,j \in V} X_{ij} = |S|$. Therefore, we get:

$$\alpha(G) \leq \vartheta(G) := \max_{X \succeq 0} \{\langle X, J \rangle : \mathrm{Tr}(X) = 1,\ X_{ij} = 0 \quad \forall \{i,j\} \in E\}.$$

Bachoc, Nebe, Oliveira, and Vallentin [**BNdOFV09**] extended ϑ to infinite geometric graphs on compact metric spaces. They also showed that this extension leads to the classical linear programming bound for spherical codes of Delsarte, Goethals, and Seidel [**DGS77**]; the linear programming bound of Cohn and Elkies for the sphere-packing density [**CE03**] can also be seen as an appropriate extension of ϑ [**dLdOFV14, dOFV19**]. These many applications illustrate the power of the Lovász theta number as a unifying concept in optimization; Goemans [**Goe97**] even remarked that "it seems all paths lead to ϑ!".

Outline. In this chapter, we show how a recursive variant of ϑ can be used to find explicit upper bounds for the independence ratio of certain geometric *hyper*graphs on the sphere and on the Euclidean space; this will lead to new bounds for a problem in Euclidean Ramsey theory. As we discuss at the end of the chapter, our recursive ϑ-number may also be applied to geometric graphs on the binary cube $\{0,1\}^n$. Analysis of the resulting upper bounds, however, is more difficult in that setting.

10.1. Overview of the construction

Let us begin by giving a high-level overview of the construction of our recursive theta number. Let $G = (V, E)$ be a (finite) graph. For $k \geq 2$, we say a subset $S \subseteq V$ contains a k-clique if there are k vertices $v_1, \ldots, v_k \in S$ such that $\{v_i, v_j\}$ is an edge for all distinct $i, j \in [k]$. That is, the complete graph K_k is a subgraph of the induced graph $G[S]$. One may then consider the parameter:

$$\alpha(G, k) := \max_{S \subseteq V} \{|S| : S \text{ contains no } k\text{-clique}\}.$$

Note that for $k = 2$, we simply have the independence number $\alpha(G, 2) = \alpha(G)$, and that $\alpha(G, 1) = 0$ by definition. The parameter $\alpha(G, k)$ can also be seen as the independence number of the (regular) hypergraph with vertex set V whose hyperedges are the k-cliques of G.

The key observation is the following. Suppose $S \subseteq V$ contains no k-clique, and let $v \in S$. Write $N_G(v) = \{w \in V : \{v, w\} \in E\}$ for the neighboorhood of v in G. Then the intersection $S \cap N_G(v)$ cannot contain a $(k-1)$-clique. Indeed, the union of such a clique with v would form a k-clique in S. Therefore, the incidence vector $\mathbf{x}_S$ of S satisfies:

$$(\mathbf{x}_S)_v \cdot \sum_{i \in N_G(v)} (\mathbf{x}_S)_i \leq \alpha(G[N_G(v)],\ k-1) \quad (v \in V). \tag{10.2}$$

The constraints (10.2) allow us to formulate a semidefinite upper bound on $\alpha(G,k)$ in terms of the parameters $\alpha(G[N_G(v)],k-1)$, which takes a similar form to ϑ. The issue is that these parameters are not known to us. Note however, that we may replace the RHS of (10.2) by any *upper bound* on $\alpha(G[N_G(v)],\ k-1)$. For instance, we could replace it by the upper bound arising from the inequalities (10.2) for $\alpha(G[N_G(v)],\ k-1)$ itself (thus in terms of parameters of the form $\alpha(\cdot,\ k-2)$). After recursively applying this procedure $k-2$ times, we end up with a bound in terms of parameters of the form $\alpha(\cdot,\ 2) = \alpha(\cdot)$, which we may upper bound using the (regular) theta number! This finally leads to a semidefinite upper bound on $\alpha(G,k)$. If $k \in \mathbb{N}$ is *fixed*, this bound may be computed in polynomial time, see [**CSdOFSV22**].

The approach outlined above may be formulated entirely in the language of (regular) hypergraphs. Furthermore, it can then be shown that some of the fundamental geometrical properties of the feasible region of the program defining the classical theta number are preserved for our recursive bound. This is the main subject of the work [**CSdOFSV22**], which we do not cover fully in this thesis. Rather, we focus here on an application of the idea presented above to geometric graphs, particularly on the hypersphere and on the Euclidean space. As we shall see below, the symmetry exhibited by these graphs allows us to formulate our recursive ϑ number in a very elegant way, permitting us to find analytical expressions for the resulting upper bounds.

10.2. Main results

10.2.1. The unit sphere. We call a set $\{\mathbf{x}^{(1)},\dots,\mathbf{x}^{(k)}\}$ of $k \geq 2$ points in the $(n-1)$-dimensional unit sphere $S^{n-1} = \{\,\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| = 1\,\}$ a (k,t)-*simplex* if $\mathbf{x}^{(i)} \cdot \mathbf{x}^{(j)} = t$ for all $i \neq j$. Note that a (k,t)-simplex has dimension $k-1$. There is a (k,t)-simplex in S^{n-1} for every $k \leq n$ and $t \in [-1/(k-1), 1)$.

Fix $n \geq k \geq 2$ and $t \in [-1/(k-1), 1)$. A set of points in S^{n-1} *avoids* (k,t)-simplices if no k points in the set form a (k,t)-simplex. We are interested in the parameter

$$\alpha(S^{n-1},k,t) = \sup_{I \subseteq S^{n-1}} \{\,\omega(I) : I \text{ is measurable and avoids } (k,t)\text{-simplices}\,\},$$

where ω is the surface measure on the sphere normalized so the total measure is 1. This is the independence ratio of the hypergraph whose vertex set is S^{n-1} and whose edges are all unit (k,t)-simplices.

In Section 10.3 we will define the parameter $\vartheta(S^{n-1},k,t)$ recursively as the optimal value of the problem

$$\begin{aligned}
\sup\ & \textstyle\int_{S^{n-1}} \int_{S^{n-1}} f(\mathbf{x}\cdot\mathbf{y})\, d\omega(\mathbf{y})d\omega(\mathbf{x}) \\
& f(1) = 1, \\
& f(t) \leq \vartheta(S^{n-2}, k-1, t/(1+t)), \\
& f \in C([-1,1]) \text{ is a function of positive type for } S^{n-1}
\end{aligned}$$

for $k \geq 3$ (see Section 10.3 for the definition of functions of positive type). The base of the recursion is $k = 2$: $\vartheta(S^{n-1}, 2, t)$ is the optimal value of the problem above when "$f(t) \leq \vartheta(S^{n-2}, k-1, t/(1+t))$" is replaced by "$f(t) = 0$".

From Theorem 10.3 below it follows that $\vartheta(S^{n-1}, k, t) \geq \alpha(S^{n-1}, k, t)$. Using extremal properties of Gegenbauer polynomials, an explicit formula can be computed for this bound, as shown in Theorem 10.5.

10.2.2. The Euclidean space. Transferring these concepts from the compact unit sphere to the non-compact Euclidean space requires a bit of care; this is done in Section 10.4.

A *unit k-simplex* in $\mathbb{R}^n$ is a set $\{\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(k)}\}$ of $k \leq n+1$ points such that $\|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\| = 1$ for all $i \neq j$. As before, note that the dimension of a unit k-simplex is $k-1$. A set of points in $\mathbb{R}^n$ *avoids* unit k-simplices if no k points in the set form a unit k-simplex. We are interested in the parameter

$$\alpha(\mathbb{R}^n, k) = \sup\{\, \overline{\delta}(I) : I \subseteq \mathbb{R}^n \text{ is measurable and avoids unit } k\text{-simplices}\,\},$$

where $\overline{\delta}(X)$ is the *upper density* of $X \subseteq \mathbb{R}^n$, that is,

$$\overline{\delta}(X) = \limsup_{T \to \infty} \frac{\operatorname{vol}(X \cap [-T, T]^n)}{\operatorname{vol}[-T, T]^n}.$$

Again, this parameter has an interpretation in terms of the independence ratio of a hypergraph on the Euclidean space and again we can bound the independence ratio from above by an appropriately defined parameter $\vartheta(\mathbb{R}^n, k)$. Theorem 10.7 below gives an explicit expression for $\vartheta(\mathbb{R}^n, k)$ in terms of Bessel functions and Gegenbauer (ultraspherical) polynomials.

The key point is that the maximum density $\alpha(\mathbb{R}^n, k)$ of unit k-simplices in $\mathbb{R}^n$ is related to the maximum density $\alpha(S^{n-1}, k-1, 1/2)$ of $(k-1, 1/2)$-simplices in the hypersphere. This is made precise in Theorem 10.3 below.

10.2.3. Euclidean Ramsey theory. The central question of Euclidean Ramsey theory is: given a finite configuration P of points in $\mathbb{R}^n$ and an integer $r \geq 1$, does every r-coloring of $\mathbb{R}^n$ contain a monochromatic congruent copy of P?

The simplest point configurations are unit k-simplices, which are known to have the exponential Ramsey property: the minimum number $\chi(\mathbb{R}^n, k)$ of colors needed to color the points of $\mathbb{R}^n$ in such a way that there are no monochromatic unit k-simplices grows exponentially in n. This was first proved by Frankl and Wilson [**FW81**] for $k = 2$ and by Frankl and Rödl [**FR87**] for $k > 2$. Results in this area are usually proved by the linear algebra method; see also Sagdeev [**Sag18a**].

Recently, Naslund [**Nas20**] used the slice-rank method from the work of Croot, Lev, and Pach [**CLP17**] and Ellenberg and Gijswijt [**EG17**] on the cap-set problem[1] to prove that

$$\chi(\mathbb{R}^n, 3) \geq (1.01466 + o(1))^n.$$

This is the best lower bound known at the moment.

For simplices of higher dimension there are also explicit lower bounds. Currently, the best such bound was obtained by Sagdeev [**Sag18b**] using a quantitative version of the Frankl-Rödl theorem:

$$\chi(\mathbb{R}^n, k) \geq \left(1 + \frac{1}{2^{2^{k+3}}} + o(1)\right)^n.$$

Denote by $H(n, k)$ the *unit-distance hypergraph*, namely the k-uniform hypergraph whose vertex set is $\mathbb{R}^n$ and whose edges are all unit k-simplices. The parameter $\chi(\mathbb{R}^n, k)$ is the chromatic number of this hypergraph. A theorem of de Bruijn and Erdős [**dBE51**] shows that computing $\chi(\mathbb{R}^n, k)$ is a combinatorial problem: the chromatic number of $H(n, k)$ is the maximum chromatic number of any finite subgraph of $H(n, k)$.

The combinatorial nature of the chromatic number makes it hard to use analytical tools in its study. This led Falconer [**Fal81**], in the case $k = 2$, to introduce the measurable counterpart of $\chi(\mathbb{R}^n, k)$, denoted by $\chi_{\mathrm{m}}(\mathbb{R}^n, k)$, by requiring the color classes to be Lebesgue-measurable sets. Of course, $\chi_{\mathrm{m}}(\mathbb{R}^n, k) \geq \chi(\mathbb{R}^n, k)$, but it is not known whether the two numbers differ.

The restriction to measurable color classes is natural and allows us to use the analytical tools developed in this chapter. Since

$$\alpha(\mathbb{R}^n, k)\chi_{\mathrm{m}}(\mathbb{R}^n, k) \geq 1,$$

any upper bound for $\alpha(\mathbb{R}^n, k)$ gives a lower bound for $\chi_{\mathrm{m}}(\mathbb{R}^n, k)$, hence

$$\chi_{\mathrm{m}}(\mathbb{R}^n, k) \geq \lceil 1/\vartheta(\mathbb{R}^n, k) \rceil.$$

In Section 10.5 we analyze the upper bounds $\vartheta(S^{n-1}, k, t)$ for simplex-avoiding sets on the sphere and $\vartheta(\mathbb{R}^n, k)$ for simplex-avoiding sets on the Euclidean space by using properties of Gegenbauer (ultraspherical) polynomials, obtaining the following theorem.

THEOREM 10.1. *If $k \geq 2$, then:*

(i) *for every $t \in (0, 1)$, there is a constant $c = c(k, t) \in (0, 1)$ such that $\vartheta(S^{n-1}, k, t) \leq (c + o(1))^n$;*
(ii) *there is a constant $c = c(k) \in (0, 1)$ such that $\vartheta(\mathbb{R}^n, k) \leq (c + o(1))^n$.*

From this theorem we get an exponential lower bound for $\chi_{\mathrm{m}}(\mathbb{R}^n, k)$. Rigorous estimates of the constant c then yield significantly better lower bounds for $\chi_{\mathrm{m}}(\mathbb{R}^n, k)$ than those coming from $\chi(\mathbb{R}^n, k)$.

---

[1]The slice-rank method is only implicit in the original works; the actual notion of slice rank for a tensor was introduced by Tao in a blog post [**Tao16**].

Indeed, in the case $k = 3$ we obtain (see Section 10.5.1)

$$\alpha(\mathbb{R}^n, 3) \le (0.95622 + o(1))^n,$$

and so

$$\chi_{\mathrm{m}}(\mathbb{R}^n, 3) \ge (1.04578 + o(1))^n.$$

We also obtain the rougher estimate

$$\alpha(\mathbb{R}^n, k) \le \left(1 - \frac{1}{9(k-1)^2} + o(1)\right)^n,$$

valid for all $k \ge 3$, which immediately implies

$$\chi_{\mathrm{m}}(\mathbb{R}^n, k) \ge \left(1 + \frac{1}{9(k-1)^2} + o(1)\right)^n.$$

Though our lower bounds for $\chi_{\mathrm{m}}(\mathbb{R}^n, k)$ do not necessarily hold for $\chi(\mathbb{R}^n, k)$, they do imply some structure for general colorings. If a coloring of $H(n, k)$ uses fewer than $1/\alpha(\mathbb{R}^n, k)$ colors, then the closure of one of the color classes is a measurable set with density greater than $\alpha(\mathbb{R}^n, k)$, and so it contains a unit k-simplex. This means that in such a coloring there are monochromatic k-point configurations arbitrarily close to unit k-simplices.

10.2.4. Notation and preliminaries. We will denote the Euclidean inner product between $\mathbf{x}$, $\mathbf{y} \in \mathbb{R}^n$ by $\mathbf{x} \cdot \mathbf{y}$. The surface measure on the sphere is denoted by ω and is always normalized so the total measure is 1.

We always normalize the (left-invariant) Haar measure on a compact group so the total measure is 1. By $\mathrm{O}(n)$ we denote the group of $n \times n$ orthogonal matrices. If $X \subseteq S^{n-1}$ is any measurable set and if μ is the Haar measure on $\mathrm{O}(n)$, then for every $e \in S^{n-1}$ we have

$$\mu(\{\, T \in \mathrm{O}(n) : Te \in X \,\}) = \omega(X).$$

We will need the following technical lemma, which will be applied to the sphere and the torus. For a proof, see Lemma 5.5 in DeCorte, Oliveira, and Vallentin [**DdOFV20**].

LEMMA 10.2. *Let V be a metric space and Γ be a compact group that acts transitively on V; let ν be a finite Borel measure on V that is positive on open sets. Denote by μ the Haar measure on Γ. If the metric on V and the measure ν are Γ-invariant and if $f \in L^2(V; \nu)$, then the function $\mathrm{K}\colon V \times V \to \mathbb{R}$ such that*

$$\mathrm{K}(\mathbf{x}, \mathbf{y}) = \int_\Gamma f(\sigma\mathbf{x}) f(\sigma\mathbf{y})\, d\mu(\sigma)$$

is continuous.

10.3. Simplex-avoiding sets on the sphere

We call a continuous kernel K: $S^{n-1} \times S^{n-1} \to \mathbb{R}$ *positive* if for every finite set $U \subseteq S^{n-1}$ the matrix $\big(\mathrm{K}(\mathbf{x}, \mathbf{y})\big)_{\mathbf{x},\mathbf{y} \in U}$ is positive semidefinite. A continuous function $f\colon [-1, 1] \to \mathbb{R}$ is of *positive type for* S^{n-1} if the kernel $\mathrm{K} \in C(S^{n-1} \times S^{n-1})$ given by $\mathrm{K}(\mathbf{x}, \mathbf{y}) = f(\mathbf{x} \cdot \mathbf{y})$ is positive.

Fix $n \geq k \geq 3$ and $t \in [-1/(k-1), 1)$. For any $\gamma \geq 0$, consider the optimization problem

$$\begin{aligned} \sup\ & \textstyle\int_{S^{n-1}} \int_{S^{n-1}} f(\mathbf{x} \cdot \mathbf{y})\, d\omega(\mathbf{y}) d\omega(\mathbf{x}) \\ & f(1) = 1, \\ & f(t) \leq \gamma, \\ & f \in C([-1, 1]) \text{ is a function of positive type for } S^{n-1}. \end{aligned} \tag{10.3}$$

The following theorem is the main technical result required to derive our recursive theta number for S^{n-1} in Theorem 10.4 below.

THEOREM 10.3. *Fix $n \geq k \geq 3$, $t \in [-1/(k-1), 1)$. If $\gamma \geq \alpha(S^{n-2}, k-1, t/(1+t))$, then the optimal value of* (10.3) *is an upper bound for $\alpha(S^{n-1}, k, t)$.*

PROOF. Let $I \subseteq S^{n-1}$ be a measurable set that avoids (k, t)-simplices and assume $\omega(I) > 0$. Consider the kernel K: $S^{n-1} \times S^{n-1} \to \mathbb{R}$ such that

$$\mathrm{K}(\mathbf{x}, \mathbf{y}) = \int_{\mathrm{O}(n)} \chi_I(T\mathbf{x}) \chi_I(T\mathbf{y})\, d\mu(T),$$

where χ_I is the characteristic function of I and where μ is the Haar measure on $\mathrm{O}(n)$.

By taking $V = S^{n-1}$ and $\Gamma = \mathrm{O}(n)$ in Lemma 10.2, we see that K is continuous. By construction, K is also positive and invariant, that is, $\mathrm{K}(T\mathbf{x}, T\mathbf{y}) = \mathrm{K}(\mathbf{x}, \mathbf{y})$ for all $T \in \mathrm{O}(n)$ and $\mathbf{x}$, $\mathbf{y} \in S^{n-1}$. Such kernels are of the form $\mathrm{K}(\mathbf{x}, \mathbf{y}) = g(\mathbf{x} \cdot \mathbf{y})$, where $g \in C([-1, 1])$ is of positive type for S^{n-1}. Note that

$$\mathrm{K}(\mathbf{x}, \mathbf{x}) = \int_{\mathrm{O}(n)} \chi_I(T\mathbf{x})\, d\mu(T) = \omega(I),$$

so $g(1) = \omega(I) > 0$.

Set $f = g/g(1)$. Immediately we have that f is continuous and of positive type and that $f(1) = 1$; moreover

$$\int_{S^{n-1}} \int_{S^{n-1}} f(\mathbf{x} \cdot \mathbf{y})\, d\omega(\mathbf{y}) d\omega(\mathbf{x}) = \omega(I).$$

Hence, if we show that $f(t) \leq \gamma$, the theorem will follow.

If $\mathbf{x} \in S^{n-1}$ is a point in a (k, t)-simplex, all other points in the simplex lie in the *link* $U_{\mathbf{x},t} = \{\, \mathbf{y} \in S^{n-1} : \mathbf{y} \cdot \mathbf{x} = t \,\}$. Note that $U_{\mathbf{x},t}$ is an $(n-2)$-dimensional sphere with radius $(1 - t^2)^{1/2}$ (see also Figure 10.1); let ν be the surface measure on $U_{\mathbf{x},t}$ normalized so the total measure is 1.

If $T \in \mathrm{O}(n)$ is any orthogonal matrix, then TI avoids (k, t)-simplices. Hence if $\mathbf{x} \in TI$, then $TI \cap U_{\mathbf{x},t}$ cannot contain $k-1$ points with pairwise

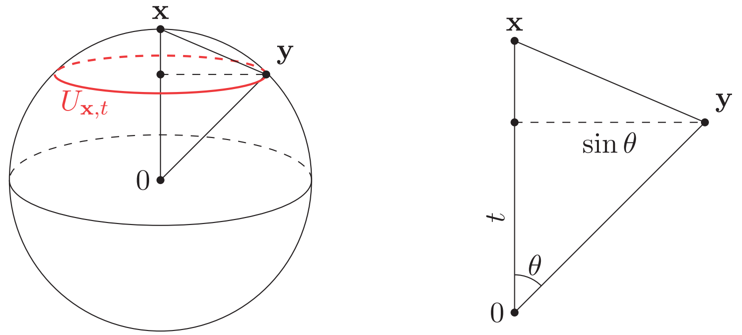

FIGURE 10.1. The link $U_{\mathbf{x},t} = \{\mathbf{y} \in S^{n-1} : \mathbf{y}\cdot\mathbf{x} = t\}$. Note that $\cos\theta = t$, implying that $U_{\mathbf{x},t}$ has radius $\sin\theta = \sqrt{1-\cos^2\theta} = \sqrt{1-t^2}$.

inner product t, and so $\nu(TI \cap U_{\mathbf{x},t}) \le \alpha(S^{n-2}, k-1, t/(1+t)) \le \gamma$. Indeed, the natural bijection between $U_{\mathbf{x},t}$ and S^{n-2} maps pairs of points with inner product t to pairs of points with inner product $t/(1+t)$, and so $TI \cap U_{\mathbf{x},t}$ is mapped to a subset of S^{n-2} avoiding $(k-1, t/(1+t))$-simplices.

Now fix $\mathbf{x} \in S^{n-1}$ and note that $g(t) = \mathrm{K}(\mathbf{x},\mathbf{y})$ for any $\mathbf{y} \in U_{\mathbf{x},t}$. An averaging argument now shows that:

$$\begin{aligned} g(t) = \int_{U_{\mathbf{x},t}} \mathrm{K}(\mathbf{x},\mathbf{y})\, d\nu(\mathbf{y}) &= \int_{U_{\mathbf{x},t}} \int_{\mathrm{O}(n)} \chi_I(T\mathbf{x})\chi_I(T\mathbf{y})\, d\mu(T) d\nu(\mathbf{y}) \\ &= \int_{\mathrm{O}(n)} \chi_I(T\mathbf{x}) \int_{U_{\mathbf{x},t}} \chi_I(T\mathbf{y})\, d\nu(\mathbf{y}) d\mu(T) \\ &\le \gamma\omega(I), \end{aligned}$$

whence $f(t) = g(t)/\omega(I) \le \gamma$, and we are done. □

One obvious choice for γ in Problem (10.3) is the bound given by the same problem for $(k-1, t/(1+t))$-simplices. The base for the recursion is $k = 2$: then we need an upper bound for the measure of a set of points on the sphere that avoids pairs of points with a fixed inner product. Such a bound was given by Bachoc, Nebe, Oliveira, and Vallentin [**BNdOFV09**] and looks very similar to (10.3). They show that, for $n \ge 2$ and $t \in [-1, 1)$, the optimal value of the following optimization problem is an upper bound for $\alpha(S^{n-1}, 2, t)$:

$$\begin{aligned} &\sup \textstyle\int_{S^{n-1}} \int_{S^{n-1}} f(\mathbf{x}\cdot\mathbf{y})\, d\omega(\mathbf{y}) d\omega(\mathbf{x}) \\ &\quad f(1) = 1, \\ &\quad f(t) = 0, \\ &\quad f \in C([-1,1]) \text{ is a function of positive type for } S^{n-1}. \end{aligned}$$

Combining this with Theorem 10.3 yields the following.

THEOREM 10.4. *Let $\vartheta(S^{n-1}, 2, t)$ denote the optimal value of the optimization problem above, so $\vartheta(S^{n-1}, 2, t) \geq \alpha(S^{n-1}, 2, t)$. For $k \geq 3$ and $t \in [-1/(k-1), 1)$, let $\vartheta(S^{n-1}, k, t)$ be the optimal value of Problem* (10.3) *when $\gamma = \vartheta(S^{n-2}, k-1, t/(1+t))$. We then have*

$$\vartheta(S^{n-1}, k, t) \geq \alpha(S^{n-1}, k, t).$$

There is actually a simple analytical expression for $\vartheta(S^{n-1}, k, t)$, as we see now. For $n \geq 2$ and $j \geq 0$, let $\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}$ denote the Gegenbauer polynomials with parameter $\alpha = (n-3)/2$ and degree j, normalized so $\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(1) = 1$, see Chapter 1.

In Theorem 6.2 of Bachoc, Nebe, Oliveira, and Vallentin [**BNdOFV09**] it is shown that for every $t \in [-1, 1)$ there is some $j \geq 0$ such that $\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t) < 0$. Theorem 8.21.8 in the book by Szegő [**Sze75**] implies that, for every $t \in (-1, 1)$,

$$\lim_{j \to \infty} \overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t) = 0.$$

Hence, for every $t \in (-1, 1)$ we can define

$$M_n(t) = \min\{\, \overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t) : j \geq 0\,\}, \tag{10.4}$$

and we see that $M_n(t) < 0$. With this we have [**BNdOFV09**, Theorem 6.2]

$$\vartheta(S^{n-1}, 2, t) = \frac{-M_n(t)}{1 - M_n(t)}.$$

The expression for $\vartheta(S^{n-1}, k, t)$ is very similar, but requires some work to derive.

THEOREM 10.5. *If $n \geq k \geq 3$ and if $t \in [-1/(k-1), 1)$, then*

$$\vartheta(S^{n-1}, k, t) = \frac{\vartheta(S^{n-2}, k-1, t/(1+t)) - M_n(t)}{1 - M_n(t)}. \tag{10.5}$$

The proof requires the following characterization of functions of positive type (see the beginning of Section 10.3 above) due to Schoenberg [**Sch42**]. A function $f\colon [-1, 1] \to \mathbb{R}$ is continuous and of positive type for S^{n-1} if and only if there are nonnegative numbers $f_0, f_1, \ldots$ such that $\sum_{j=0}^{\infty} f_j < \infty$ and

$$f(t) = \sum_{j=0}^{\infty} f_j \overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t),$$

with uniform convergence in $[-1, 1]$. Compare to the summation formula (1.15).

PROOF OF THEOREM 10.5. Using the orthogonality of Gegenbauer polynomials (see also (1.15)), we see that for any $i \neq j$:

$$\int_{S^{n-1}} \int_{S^{n-1}} \overline{\mathcal{G}}_i^{(\frac{n-3}{2})}(\mathbf{x} \cdot \mathbf{y}) \overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(\mathbf{x} \cdot \mathbf{y})\, d\omega(\mathbf{y}) d\omega(\mathbf{x}) = 0$$

As $\overline{\mathcal{G}}_0^{(\frac{n-3}{2})} = 1$, this implies in particular that, if $j \geq 1$, we have:

$$\int_{S^{n-1}} \int_{S^{n-1}} \overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(\mathbf{x} \cdot \mathbf{y})\, d\omega(\mathbf{y}) d\omega(\mathbf{x}) = 0.$$

We may use this and Schoenberg's characterization of positive type functions to rewrite (10.3) with $\gamma = \vartheta(S^{n-2}, k-1, t/(1+t))$, obtaining the equivalent problem:

$$\begin{aligned} \sup\ & f_0 \\ & \textstyle\sum_{j=0}^{\infty} f_j = 1, \\ & \textstyle\sum_{j=0}^{\infty} f_j \overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t) \leq \vartheta(S^{n-2}, k-1, t/(1+t)), \\ & f_j \geq 0 \text{ for all } j \geq 0. \end{aligned}$$

It remains to solve this problem. Note that

$$\sum_{j=0}^{\infty} f_j \overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t)$$

is a convex combination of the numbers $\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t)$. We want to keep this convex combination below $\vartheta(S^{n-2}, k-1, t/(1+t))$ while maximizing f_0. The best way to do so is to concentrate all the weight of the combination on f_0 and f_{j^*}, where j^* is such that $\overline{\mathcal{G}}_{j^*}^{(\frac{n-3}{2})}(t)$ is the most negative number appearing in the convex combination, meaning $\overline{\mathcal{G}}_{j^*}^{(\frac{n-3}{2})}(t) = M_n(t)$.

That is, we may assume that the only nonzero variables in an optimum solution are f_0 and f_{j^*}. It thus remains to maximize f_0 under the conditions that $f_0 + f_{j^*} = 1$, and:

$$f_0 + f_{j^*} \overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t) = f_0 + f_{j^*} M_n(t) \leq \vartheta(S^{n-2}, k-1, t/(1+t)),$$

which precisely yields the optimal value in the statement of the theorem. □

The expression for $\vartheta(S^{n-1}, k, 0)$ is particularly simple. Indeed, for $n \geq 2$ it follows from the recurrence relation for the Jacobi polynomials that $M_n(0) = \overline{\mathcal{G}}_2^{(\frac{n-3}{2})}(0) = -1/(n-1)$, whence

$$\vartheta(S^{n-1}, k, 0) = (k-1)/n.$$

Figure 10.2 shows the behavior of $\vartheta(S^{n-1}, 3, t)$ for a few values of n as t changes. Plots for $k > 3$ are very similar.

10.4. Simplex-avoiding sets in Euclidean space

An optimization problem similar to (10.3) provides an upper bound for $\alpha(\mathbb{R}^n, k)$. To introduce it, we need some definitions and facts from harmonic analysis on $\mathbb{R}^n$; for background, see e.g. the book by Reed and Simon [**RS75**].

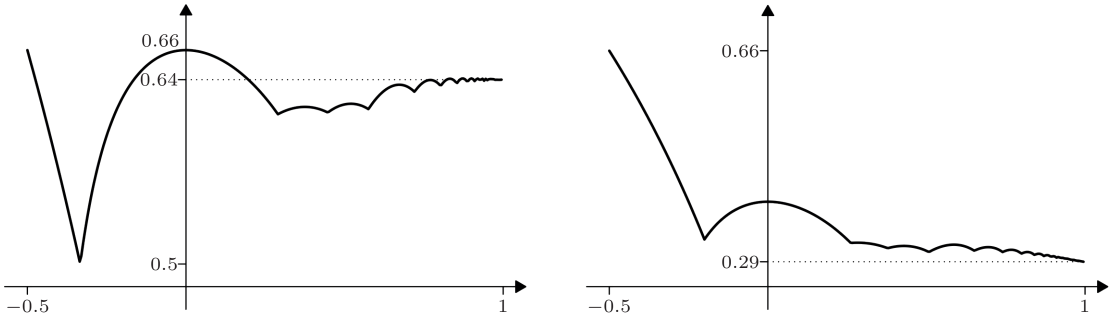

FIGURE 10.2. Plots of $\vartheta(S^{n-1}, 3, t)$ for $t \in [-0.5, 1]$ and $n = 3$ (left) and 5 (right).

A continuous function $f\colon \mathbb{R}^n \to \mathbb{R}$ is of *positive type* if for every finite set $U \subseteq \mathbb{R}^n$ the matrix $\bigl(f(\mathbf{x}-\mathbf{y})\bigr)_{\mathbf{x},\mathbf{y}\in U}$ is positive semidefinite. Such a function f has a well-defined *mean value*

$$M(f) = \lim_{T\to\infty} \frac{1}{\operatorname{vol}[-T,T]^n} \int_{[-T,T]^n} f(\mathbf{x})\, d\mathbf{x}.$$

We say that a function $f\colon \mathbb{R}^n \to \mathbb{R}$ is *radial* if $f(\mathbf{x})$ depends only on $\|\mathbf{x}\|$. In this case, for $t \geq 0$ we denote by $f(t)$ the common value of f for vectors of norm t.

Fix $n \geq 2$ and $k \geq 3$ such that $k \leq n+1$. For every $\gamma \geq 0$, consider the optimization problem

$$\begin{aligned} \sup\ & M(f) \\ & f(0) = 1, \\ & f(1) \leq \gamma, \\ & f\colon \mathbb{R}^n \to \mathbb{R} \text{ is continuous, radial, and of positive type.} \end{aligned} \tag{10.6}$$

We have the analogue of Theorem 10.3:

THEOREM 10.6. *Fix $n \geq 2$ and $k \geq 3$ such that $k \leq n+1$. If $\gamma \geq \alpha(S^{n-1}, k-1, 1/2)$, then the optimal value of* (10.6) *is an upper bound for $\alpha(\mathbb{R}^n, k)$.*

We need a few facts about periodic sets and functions. A set $X \subseteq \mathbb{R}^n$ is *periodic* if it is invariant under some lattice Λ, that is, if $X + v = X$ for all $v \in \Lambda$. Similarly, a function $f\colon \mathbb{R}^n \to \mathbb{R}$ is *periodic* if there is a lattice Λ such that $f(\mathbf{x}+v) = f(\mathbf{x})$ for all $v \in \Lambda$. We say that Λ is a *periodicity lattice* of X or f. A periodic function f with periodicity lattice Λ can be seen as a function on the torus $\mathbb{R}^n/\Lambda$; its mean value is

$$\frac{1}{\operatorname{vol}(\mathbb{R}^n/\Lambda)} \int_{\mathbb{R}^n/\Lambda} f(\mathbf{x})\, d\mathbf{x}.$$

Proof of Theorem 10.6. Let $I \subseteq \mathbb{R}^n$ be a measurable set of positive upper density avoiding unit k-simplices. The first step is to see that we can assume that I is periodic. Indeed, fix $R > 1/2$. Erase a border of width $1/2$ around $I \cap [-R, R]^n$ and paste the resulting set periodically in such a way that there is an empty gap of width 1 between any two pasted copies. The resulting periodic set still avoids unit k-simplices and is measurable. Its upper density is

$$\frac{\operatorname{vol}(I \cap [-R + 1/2, R - 1/2]^n)}{\operatorname{vol}[-R, R]^n};$$

by taking R large enough, we can make this density as close as we want to the upper density of I.

Assume then that I is periodic, so its characteristic function χ_I is also periodic; let Λ be a periodicity lattice of I. Set

$$g(\mathbf{x}) = \frac{1}{\operatorname{vol}(\mathbb{R}^n/\Lambda)} \int_{\mathbb{R}^n/\Lambda} \chi_I(\mathbf{y})\chi_I(\mathbf{x} + \mathbf{y})\, d\mathbf{y}.$$

Lemma 10.2 with $V = \Gamma = \mathbb{R}^n/\Lambda$ applied to χ_I implies that g is continuous. Direct verification yields that g is of positive type, $g(0) = \overline{\delta}(I)$, and $M(g) = \overline{\delta}(I)^2$.

Now set

$$f(\mathbf{x}) = \overline{\delta}(I)^{-1} \int_{\mathrm{O}(n)} g(T\mathbf{x})\, d\mu(T),$$

where μ is the Haar measure on $\mathrm{O}(n)$. Note that f is continuous, radial, and of positive type. Moreover, $f(0) = 1$ and $M(f) = \overline{\delta}(I)$. If we show that $f(1) \leq \gamma$, then f is a feasible solution of (10.6) with $M(f) = \overline{\delta}(I)$, and so the theorem will follow.

To see that $f(1) \leq \gamma$, note that if $\mathbf{x}$ is a point of a unit k-simplex in $\mathbb{R}^n$, then all the other points in the simplex lie on the unit sphere $\mathbf{x} + S^{n-1}$ centered at $\mathbf{x}$. Hence if $\mathbf{x} \in I$, then $I \cap (\mathbf{x} + S^{n-1})$ is a measurable subset of $\mathbf{x} + S^{n-1}$ that avoids $(k-1, 1/2)$-simplices, and so the measure of $I \cap (\mathbf{x} + S^{n-1})$ as a subset of the unit sphere is at most $\alpha(S^{n-1}, k-1, 1/2)$. Hence if $\xi \in \mathbb{R}^n$ is any unit vector, then

$$\begin{aligned}
f(1) &= \overline{\delta}(I)^{-1} \int_{\mathrm{O}(n)} g(T\xi)\, d\mu(T) \\
&= \overline{\delta}(I)^{-1} \int_{\mathrm{O}(n)} \frac{1}{\operatorname{vol}(\mathbb{R}^n/\Lambda)} \int_{\mathbb{R}^n/\Lambda} \chi_I(\mathbf{x})\chi_I(T\xi + \mathbf{x})\, d\mathbf{x} d\mu(T) \\
&= \overline{\delta}(I)^{-1} \frac{1}{\operatorname{vol}(\mathbb{R}^n/\Lambda)} \int_{\mathbb{R}^n/\Lambda} \chi_I(\mathbf{x}) \int_{\mathrm{O}(n)} \chi_I(T\xi + \mathbf{x})\, d\mu(T) d\mathbf{x} \\
&\leq \alpha(S^{n-1}, k-1, 1/2) \leq \gamma,
\end{aligned}$$

as we wanted. □

Denote by $\vartheta(\mathbb{R}^n, k)$ the optimal value of (10.6) when setting $\gamma = \vartheta(S^{n-1}, k-1, 1/2)$. Then $\vartheta(\mathbb{R}^n, k) \geq \alpha(\mathbb{R}^n, k)$.

An expression akin to the one for $\vartheta(S^{n-1}, k, t)$ can be derived for $\vartheta(\mathbb{R}^n, k)$. For $n \geq 2$ and $u \geq 0$, let

$$\Omega_n(u) = \Gamma(n/2)(2/u)^{(n-2)/2} J_{(n-2)/2}(u),$$

where J_α is the *Bessel function* of the first kind with parameter α. Let m_n be the global minimum of Ω_n, which is a negative number (cf. Oliveira and Vallentin [**dOFV10**]). The following theorem is the analogue of Theorem 10.5.

THEOREM 10.7. *For $n \geq 2$ we have*

$$\vartheta(\mathbb{R}^n, k) = \frac{\vartheta(S^{n-1}, k-1, 1/2) - m_n}{1 - m_n}.$$

The proof uses again a theorem of Schoenberg [**Sch38**], that this time characterizes radial and continuous functions of positive type on $\mathbb{R}^n$: these are the functions $f\colon \mathbb{R}^n \to \mathbb{R}$ such that

$$f(\mathbf{x}) = \int_0^\infty \Omega_n(z\|\mathbf{x}\|)\, d\nu(z) \tag{10.7}$$

for some finite Borel measure ν.

PROOF. If f is given as in (10.7), then $M(f) = \nu(\{0\})$ (see e.g. Section 6.2 in DeCorte, Oliveira, and Vallentin [**DdOFV20**]). Using Schoenberg's theorem, we can rewrite (10.6) (with $\gamma = \vartheta(S^{n-1}, k-1, 1/2)$) equivalently as:

$$\begin{aligned} \sup\ & \nu(\{0\}) \\ & \nu([0,\infty)) = 1, \\ & \textstyle\int_0^\infty \Omega_n(z)\, d\nu(z) \leq \vartheta(S^{n-1}, k-1, 1/2), \\ & \nu \text{ is a Borel measure.} \end{aligned}$$

We are now in the same situation as in the proof of Theorem 10.5. If z^* is such that $m_n = \Omega_n(z^*)$, then the optimal ν is supported at 0 and z^*. Solving the resulting system yields the theorem. □

Table 10.1 contains some values for $\vartheta(\mathbb{R}^n, k)$.

10.5. Exponential density decay

In this section we analyze the asymptotic behavior of $\vartheta(S^{n-1}, k, t)$ and $\vartheta(\mathbb{R}^n, k)$ as functions of n, proving Theorem 10.1.

The main step in our analysis is to understand the asymptotic behavior of

$$M_n(t) = \min\{\, \overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t) : j \geq 0 \,\},$$

as defined in (10.4). For $t \in [-1, 0)$ we have $M_n(t) \leq \overline{\mathcal{G}}_1^{(\frac{n-3}{2})}(t) = t$, and so $M_n(t)$ does not approach 0. We have seen in Section 10.3 that $M_n(0) = -1/(n-1)$, so for $t = 0$ we have that $M_n(t)$ approaches 0 linearly fast as n

n / k	3	4	5	6	7	8	9	10	11
2	0.64355	—	—	—	—	—	—	—	—
3	0.42849	0.69138	—	—	—	—	—	—	—
4	0.29346	0.49798	0.73225	—	—	—	—	—	—
5	0.20374	0.36768	0.55035	0.76580	—	—	—	—	—
6	0.15225	0.28471	0.42777	0.60262	0.79563	—	—	—	—
7	0.11866	0.22740	0.34071	0.48493	0.64681	0.81972	—	—	—
8	0.09339	0.18405	0.27471	0.39559	0.53374	0.68268	0.83882	—	—
9	0.07387	0.15030	0.22864	0.33042	0.44903	0.57816	0.71431	0.85537	—
10	0.05846	0.12340	0.19194	0.27851	0.38158	0.49496	0.61521	0.74026	0.86882

TABLE 10.1. The bound $\vartheta(\mathbb{R}^n, k)$ for $n = 2, \ldots, 10$ and $k = 3, \ldots, 11$, with values of n on each row and of k on each column.

grows. Things get interesting when $t \in (0, 1)$: then $M_n(t)$ approaches 0 exponentially fast as n grows.

THEOREM 10.8. *For every $t \in (0, 1)$ there is $c \in (0, 1)$ such that*

$$|M_n(t)| \leq (c + o(1))^n.$$

We will need the following lemma showing that, for every $t \in (0, 1)$, if $j = \Omega(n)$, then $|\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t)|$ decays exponentially in n. The statement of the lemma is quite a bit more precise than that, since we later want to do a more detailed analysis of the base of the exponential. The proof is a refinement of the analysis carried out by Schoenberg [**Sch42**].

LEMMA 10.9. *If for $\theta \in (0, \pi)$ and $\delta \in (0, \pi/2)$ we write*

$$C = (\cos^2 \theta + \sin^2 \theta \sin^2 \delta)^{1/2},$$

then $|\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(\cos \theta)| \leq \pi n^{1/2} \cos^{n-3} \delta + C^j$ for all $n \geq 3$.

PROOF. An integral representation for the ultraspherical polynomials due to Gegenbauer (take $\lambda = (n-2)/2$ in Theorem 6.7.4 from Andrews, Askey, and Roy [**AAR99**]) gives us the formula

$$\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(\cos \theta) = R(n)^{-1} \int_0^\pi F(\phi)^j \sin^{n-3} \phi \, d\phi,$$

where

$$F(\phi) = \cos \theta + i \sin \theta \cos \phi \qquad \text{and} \qquad R(n) = \int_0^\pi \sin^{n-3} \phi \, d\phi.$$

Note that $|F(\phi)|^2 = \cos^2 \theta + \sin^2 \theta \cos^2 \phi$ and that $|F(\phi)| \leq 1$. Split the integration domain into the intervals $[0, \pi/2 - \delta]$, $[\pi/2 - \delta, \pi/2 + \delta]$, and $[\pi/2 + \delta, \pi]$

to obtain

$$
\begin{aligned}
|\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(\cos\theta)| &\le R(n)^{-1}\int_0^{\pi} |F(\phi)|^j \sin^{n-3}\phi\, d\phi \\
&\le 2R(n)^{-1}\int_0^{\pi/2-\delta} \sin^{n-3}\phi\, d\phi \\
&\quad + R(n)^{-1}\int_{\pi/2-\delta}^{\pi/2+\delta} |F(\phi)|^j \sin^{n-3}\phi\, d\phi.
\end{aligned}
$$

For the first term above, note that

$$R(n) = \frac{\pi^{1/2}\Gamma(n/2-1)}{\Gamma((n-1)/2)}.$$

Take $x = (n-2)/2$ and $a = 1/2$ in (7) of Wendel [**Wen48**] to get

$$R(n)^{-1} \le \pi^{-1/2}((n-2)/2)^{1/2} < n^{1/2}.$$

Now

$$
\begin{aligned}
2R(n)^{-1}\int_0^{\pi/2-\delta} \sin^{n-3}\phi\, d\phi &\le 2n^{1/2}\int_0^{\pi/2-\delta} \sin^{n-3}(\pi/2-\delta)\, d\phi \\
&= 2n^{1/2}(\pi/2-\delta)\cos^{n-3}\delta \\
&\le \pi n^{1/2}\cos^{n-3}\delta.
\end{aligned}
$$

For the second term we get directly

$$R(n)^{-1}\int_{\pi/2-\delta}^{\pi/2+\delta} |F(\phi)|^j \sin^{n-3}\phi\, d\phi \le R(n)^{-1}\int_{\pi/2-\delta}^{\pi/2+\delta} C^j \sin^{n-3}\phi\, d\phi \le C^j,$$

and we are done. □

We can now prove the theorem.

Proof of Theorem 10.8. Our strategy is to find a lower bound on the largest j_0 such that $\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t) \ge 0$ for all $j \le j_0$. Then we know that $M_n(t)$ is attained by some $j \ge j_0$, and we can use Lemma 10.9 to estimate $|M_n(t)|$.

Recall [**Sze75**] that the zeros of $\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}$ are all in $[-1,1]$ and that the rightmost zero of $\overline{\mathcal{G}}_{j+1}^{(\frac{n-3}{2})}$ is to the right of the rightmost zero of $\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}$. For convenience, let C_j^{λ} denote the Gegenbauer polynomial with shifted parameter $\lambda - 1/2$ and degree j, so

$$\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t) = \frac{C_j^{(n-2)/2}(t)}{C_j^{(n-2)/2}(1)}. \tag{10.8}$$

Let x_j be the largest zero of C_j^λ. Elbert and Laforgia [**EL90**, p. 94] show that, for $\lambda \geq 0$,

$$x_j^2 < \frac{j^2 + 2\lambda j}{(j+\lambda)^2}.$$

If for a given j we have that

$$\frac{j^2 + 2\lambda j}{(j+\lambda)^2} \leq t^2, \tag{10.9}$$

then we know that the rightmost zero of C_j^λ is to the left of t, and so $C_j^\lambda(t) \geq 0$.

Note that the left-hand side in (10.9) is increasing in j. Let us estimate the largest j for which (10.9) holds. We want

$$j^2 + 2\lambda j - t^2(j+\lambda)^2 \leq 0.$$

The left-hand side above is quadratic in j, and since $t^2 < 1$ the coefficient of j^2 is positive. So all we have to do is to compute the largest root of the left-hand side, which is $2a(t)\lambda$, where $a(t) = ((1-t^2)^{-1/2} - 1)/2$.

Hence for $j \leq 2a(t)\lambda$ we have $C_j^\lambda(t) \geq 0$. From (10.8) we see that $\overline{\mathcal{G}}_j^{(\frac{n-3}{2})}(t) \geq 0$ if

$$j \leq a(t)n - 2a(t).$$

Now plug the right-hand side above into the upper bound of Lemma 10.9 to get

$$\begin{aligned} |M_n(t)| &\leq (\pi n^{1/2} \cos^{-3}\delta)\cos^n\delta + C^{a(t)n-2a(t)} \\ &= O(n^{1/2})\cos^n\delta + O(1)(C^{a(t)})^n, \end{aligned}$$

with C as defined in Lemma 10.9 with $\cos\theta = t$. For any choice of $\delta \in (0, \pi/2)$, we have that $\cos\delta$, $C \in (0,1)$, and since $a(t) > 0$ for all $t \in (0,1)$, the theorem follows.

□

Proof of exponential decay. We now get exponential decay for $\vartheta(S^{n-1}, k, t)$ for any $k \geq 3$ and $t \in (0,1)$. Indeed, consider the recurrence $F_0 = t$ and $F_i = F_{i-1}/(1+F_{i-1})$ for $i \geq 1$, whose solution is $F_i = t/(1+it)$. Using Theorem 10.8 to develop our analytic solution (10.5), we get

$$\vartheta(S^{n-1}, k, t) \sim \sum_{i=0}^{k-2} |M_{n-i}(F_i)| = \sum_{i=0}^{k-2} |M_{n-i}(t/(1+it))|,$$

where $a_n \sim b_n$ means that $\lim_{n\to\infty} a_n/b_n = 1$. Since $t/(1+it) > 0$ for all i, each term decays exponentially fast, and so we get exponential decay for the sum.

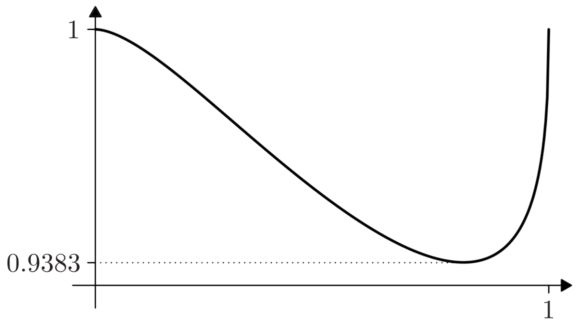

FIGURE 10.3. The best constant c obtained in our proof of Theorem 10.8, for each value of $t \in (0, 1)$.

We also get exponential decay for $\vartheta(\mathbb{R}^n, k)$ for any $k \geq 3$. Indeed, from Theorem 10.7 we have that

$$\vartheta(\mathbb{R}^n, k) \sim |m_n| + \sum_{i=0}^{k-3} |M_{n-i}(1/(2+i))|.$$

From Theorem 10.8 we know that every term in the summation above decays exponentially fast. Bachoc, Nebe, Oliveira, and Vallentin [**BNdOFV09**] give an asymptotic bound for $|m_n|$ that shows that it also decays exponentially in n, namely

$$|m_n| \leq (2/e + o(1))^{n/2} = (0.8577\ldots + o(1))^n.$$

This finishes the proof of Theorem 10.1.

10.5.1. Explicit bounds. We now compute explicit constants $c(k, t)$ and $c(k)$ which can serve as bases for the exponentials in Theorem 10.1, in particular obtaining the bounds advertised in Section 10.2.3.

The constant c given in Theorem 10.8 depends on t. Following the proof, we can find the best constant for every $t \in (0, 1)$ by finding $\delta \in (0, \pi/2)$ such that $\cos\delta = C^{a(t)}$, that is, by solving the equation

$$\cos^4 \delta = (t^2 + (1 - t^2)\sin^2 \delta)^{(1-t^2)^{-1/2}-1} \tag{10.10}$$

and taking $c = \cos\delta > 0$.

For any given $t \in (0, 1)$ it is easy to solve (10.10) numerically. For $t = 1/2$ we get $\cos\delta = 0.95621\ldots$ as a solution, and so $|M_n(1/2)| \leq (0.95622 + o(1))^n$, leading to the the bound

$$\vartheta(\mathbb{R}^n, 3) \sim |M_n(1/2)| \leq (0.95622 + o(1))^n.$$

Figure 10.3 shows a plot of the best constant c for every $t \in (0, 1)$.

With a little extra work, it is possible to show that, for all $k \geq 2$,

$$|M_n(1/k)| \leq \left(1 - \frac{1}{9k^2} + o(1)\right)^n, \tag{10.11}$$

whence

$$\vartheta(\mathbb{R}^n, k) \sim |M_{n-k+3}(1/(k-1))| \leq \left(1 - \frac{1}{9(k-1)^2} + o(1)\right)^n$$

for all $k \geq 3$.

Direct verification shows that (10.11) holds for $k = 2$, so let us assume $k \geq 3$. Writing c for the (unique) positive solution $\cos\delta$ of (10.10) and taking $\theta \in (0, \pi/2)$ such that $\cos\theta = t$, we can rewrite (10.10) in the more convenient form

$$c^{4\sin\theta/(1-\sin\theta)} = 1 - c^2 \sin^2\theta.$$

Now say $c = 1 - x$ and use Bernoulli's inequality $(1+z)^r \geq 1 + rz$ to get

$$\begin{aligned}(1-x)^{4\sin\theta/(1-\sin\theta)} &\geq 1 - \frac{4\sin\theta}{1-\sin\theta}x \quad \text{and}\\ 1-(1-x)^2\sin^2\theta &\leq 1-(1-2x)\sin^2\theta.\end{aligned}$$

Equating the left-hand sides of both inequalities above and solving for x, we get

$$c = 1 - x \leq 1 - \frac{\sin\theta(1-\sin\theta)}{4 + 2\sin\theta(1-\sin\theta)}.$$

In particular, when $\cos\theta = 1/k$ we get

$$\begin{aligned}|M_n(1/k)| &\leq \left(1 - \frac{1}{4k^2(1+\sqrt{k^2/(k^2-1)})+2} + o(1)\right)^n\\ &\leq \left(1 - \frac{1}{9k^2} + o(1)\right)^n\end{aligned}$$

for all $k \geq 3$.

10.6. Triangle-avoiding sets in the binary cube

For an integer $n \geq 1$, consider the binary cube $\mathbb{B}^n = \{0,1\}^n$ equipped with the *Hamming distance*, which for $\mathbf{x}$, $\mathbf{y} \in \mathbb{B}^n$ is denoted by $d(\mathbf{x}, \mathbf{y})$ and equals the number of bits in which $\mathbf{x}$ and $\mathbf{y}$ differ. A classical problem in coding theory is to determine the parameter $A(n,d)$, which is the maximum size of a subset I of $\mathbb{B}^n$ such that $d(\mathbf{x}, \mathbf{y}) \geq d$ for all distinct $\mathbf{x}$, $\mathbf{y} \in I$.

If we let $G(n,d)$ be the graph with vertex set $\mathbb{B}^n$ in which $\mathbf{x}$, $\mathbf{y} \in \mathbb{B}^n$ are adjacent if $d(\mathbf{x}, \mathbf{y}) < d$, then $A(n,d) = \alpha(G(n,d))$. A simple variant of the Lovász theta number of $G(n,d)$ then provides an upper bound for $A(n,d)$, which is easy to compute given the abundant symmetry of $G(n,d)$. This bound, known as the linear programming bound, was originally described by Delsarte [**Del73**]; its relation to the theta number was later discovered by McEliece, Rodemich, and Rumsey [**MRR78**] and Schrijver [**Sch79**].

Now let $s \geq 1$ be an integer. Three distinct points $\mathbf{x}^{(1)}$, $\mathbf{x}^{(2)}$, $\mathbf{x}^{(3)} \in \mathbb{B}^n$ form an *s-triangle* if $d(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = s$ for all $i \neq j$. It is easy to show that there is an s-triangle if and only if s is even and $0 < s \leq \lfloor 2n/3 \rfloor$. We want to find the largest size $\alpha(\mathbb{B}^n, 3, s)$ of a set of points in $\mathbb{B}^n$ that avoids s-triangles. This is thus the analogue of the problem we have so far considered on S^{n-1}, but restricted to 3-simplices (i.e., triangles).

We sketch how to obtain the corresponding recursive theta number; the idea of the construction is exactly the same as for S^{n-1}. See also [**CSdOFSV22**].

The functions of positive type on $\mathbb{B}^n$ are of the form:

$$f(t) = \sum_{k=0}^{n} a_k \overline{\mathcal{K}}_k^{(n)}(t),$$

where $\overline{\mathcal{K}}_k^{(n)}$ is the Krawtchouck polynomial of degree k and $a_0, \ldots, a_n \geq 0$. Compare to the summation formula (1.17). In light of this characterization, we get the following upper bound on $\alpha(\mathbb{B}^n, 3, s)$:

$$\begin{aligned} \vartheta(\mathbb{B}^n, 3, s) := \max\ & 2^n a_0 \\ & \textstyle\sum_{k=0}^n a_k = 1, \\ & \textstyle\sum_{k=0}^n a_k \overline{\mathcal{K}}_k^{(n)}(s) \leq |\mathbb{B}_s^n|^{-1} \vartheta(\mathbb{B}_s^n), \\ & a_0, \ldots, a_n \geq 0. \end{aligned} \tag{10.12}$$

Here, $\mathbb{B}_s^n$ is the the link of $0 \in \mathbb{B}^n$, which consists of the set of all words of weight s (the *weight* of a word is the number of 1s in it); two words are adjacent in $\mathbb{B}_s^n$ if they are at distance s and $\vartheta(\mathbb{B}_s^n)$ is the regular theta number. As before, the program (10.12) admits a simple analytical optimum solution.

THEOREM 10.10. *Write $M_K^n(s) = \min\{ K_k^n(s) : k = 0, \ldots, n \}$ for $s \geq 0$. If $n \geq 1$ is an integer and $0 < s \leq \lfloor 2n/3 \rfloor$ is an even integer, then we have:*

$$\alpha(\mathbb{B}^n, 3, s) \leq \vartheta(\mathbb{B}^n, 3, s) = 2^n \frac{M_K^n(s) - |\mathbb{B}_s^n|^{-1}\vartheta(\mathbb{B}_s^n)}{M_K^n(s) - 1}.$$

To compute $\vartheta(\mathbb{B}_s^n)$ we again use symmetry. Let $A \colon \mathbb{B}_s^n \times \mathbb{B}_s^n \to \mathbb{R}$ be a matrix. If A is $\mathrm{Iso}(\mathbb{B}_s^n)$-invariant, then $A(\mathbf{x}, \mathbf{y})$ depends only on $d(\mathbf{x}, \mathbf{y})$, and so we write $A(t)$ for the value of $A(\mathbf{x}, \mathbf{y})$ when $d(\mathbf{x}, \mathbf{y}) = t$. The matrix A is $\mathrm{Iso}(\mathbb{B}_s^n)$-invariant and positive semidefinite if and only if there are numbers a_0, $\ldots$, $a_s \geq 0$ such that

$$A(t) = \sum_{k=0}^{s} a_k Q_k^{n,s}(t/2)$$

(note that Hamming distances in $\mathbb{B}_s^n$ are always even), where $Q_k^{n,s}$ is the Hahn polynomial of degree k, which for an even integer $t \geq 0$ is such that

$$Q_k^{n,s}(t) = \sum_{i=0}^{k} (-1)^i \binom{s}{i}^{-1} \binom{n-s}{i}^{-1} \binom{k}{i} \binom{n+1-k}{i} \binom{t}{i}.$$

The polynomials are normalized so $Q_k^{n,s}(0) = 1$; if $E_k(\mathbf{x}, \mathbf{y}) = Q_k^{n,s}(d(\mathbf{x}, \mathbf{y})/2)$, then $\langle E_k, E_l \rangle = 0$ whenever $k \neq l$ (see Delsarte [**Del78**], in particular Theorem 5, and Dunkl [**Dun78**]).

With this characterization, $\langle J, A\rangle = |\mathbb{B}_s^n|^2 a_0$ since $E_0 = J$. We see that $\vartheta(\mathbb{B}_s^n)$ is the optimal value of the problem

$$\begin{aligned} \max\ & |\mathbb{B}_s^n| a_0 \\ & \textstyle\sum_{k=0}^s a_k = 1, \\ & \textstyle\sum_{k=0}^s a_k Q_k^{n,s}(s/2) = 0, \\ & a_0, \ldots, a_s \geq 0. \end{aligned}$$

With $M_Q^n(s) = \min\{\, Q_k^{n,s}(s/2) : k = 0, \ldots, s\,\}$, we have the following analogue of Theorem 10.10.

Theorem 10.11. *If $n \geq 1$ is an integer and $0 < s \leq \lfloor 2n/3 \rfloor$ is an even integer, then*

$$\vartheta(\mathbb{B}_s^n) = |\mathbb{B}_s^n| \frac{M_Q^n(s)}{M_Q^n(s) - 1}.$$

Putting things together, we obtain an expression for $\vartheta(\mathbb{B}^n, 3, s)$ in terms of $M_K^n(s)$ and $M_Q^n(s)$ (similar to the one we derived earlier for simplex-avoiding sets on S^{n-1}).

Using this expression, one could attempt to analyze how the density of a subset of $\mathbb{B}^n$ that avoids s-triangles behaves as n goes to infinity. For a fixed s, this question is not interesting, as $|\mathbb{B}_s^n|$ is exponentially smaller than $|\mathbb{B}^n|$. Therefore, we should consider a regime where s tends to infinity as well. For instance, what if we take as s the even integer closest to n/c for some $c > 1$? Numerical evidence (Figure 10.4) suggests that $\vartheta(\mathbb{B}^n, 3, s)/2^n$ decays exponentially fast in n for all $c > 2$; for $c = 2$, there seems to be no exponential decay.

The open problem, which we leave for future work, is whether an analytical proof for this exponential decay may be given similar to the one presented above for S^{n-1}. It has so far proven difficult to analyze the parameters $M_K^n(s)$ and $M_Q^n(s)$. It does not seem possible to transport the techniques we used for the Gegenbauer polynomials directly. In particular, we lack a suitable integral representation for the Krawtchouk and (especially) the Hahn polynomials.

10.7. Discussion

We have defined a recursive theta number for (geometric) hypergraphs, which allows us to bound the size of simplex-avoiding sets on the hypersphere S^{n-1}. We express the resulting bound in terms of a parameter related to Gegenbauer polynomials. After careful analysis of this parameter, we are able to show exponential decay of the density of simplex-avoiding sets as n tends to infinity. This in turn allows us to show an improved version of an existing result in Euclidean Ramsey theory. Furthermore, we have seen that our approach may be applied to obtain bounds on the size of triangle-free sets in the binary cube, which may similarly by expressed in terms of parameters related now to Krawtchouk and Hahn polynomials. These parameters, unfortunately, appear to be much harder to analyze.

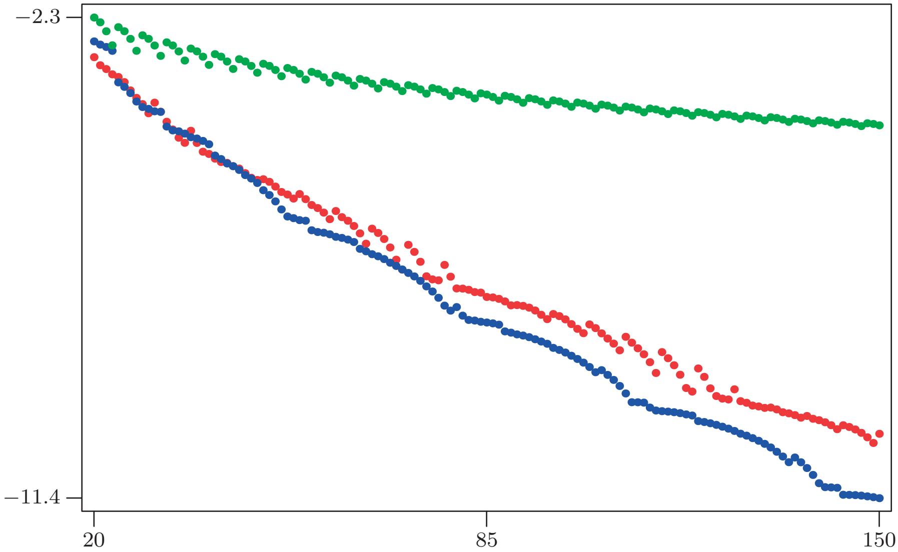

FIGURE 10.4. The plot shows, for every $n = 20, \ldots, 150$ on the horizontal axis, the value of $\ln(\vartheta(\mathbb{B}^n, 3, s)/2^n)$ on the vertical axis, where s is the even integer closest to $n/2$ (in green), $n/3$ (in red), and $n/4$ (in blue).

A recursive theta body for hypergraphs. The theta body of a graph is a convex relaxation of the independent-set polytope, defined by Grötschel, Lovász, and Schrijver [**GLS88**]. As the name suggest, it is closely related to the theta number. Indeed, ϑ is obtained by optimizing a linear function over the theta body. In the work [**CSdOFSV22**], we recursively extend the theta body to uniform hypergraphs, using ideas similar to the ones presented in this chapter. We show that some of the nice structural properties of the regular theta body are preserved in this setting.

Relation to other bounds. Finally, let us compare our recursive theta number to some similar, existing bounds in the literature. First, there is the classical *Hoffman bound* $h(G)$, which bounds the stability number $\alpha(G)$ of a regular graph G in terms of the smallest eigenvalue of its adjacency matrix. The theta number is known to be a tighter bound on $\alpha(G)$ than the Hoffman bound: $\alpha(G) \leq \vartheta(G) \leq h(G)$ for regular graphs G [**Lov79**]. Filmus, Golubev, and Lifshitz [**FGL21**] extend the Hoffman bound to the setting of (edge-weighted) hypergraphs, and give applications in extremal combinatorics. As we show in [**CSdOFSV22**], our resursive theta number is always at least as tight as their high-dimensional Hoffman bound.

Second, there are the semidefinite hierarchies of so-called *k-point bounds*, which may be thought of as higher levels of a Lasserre-type hierarchy of bounds on $\alpha(G)$ (just as ϑ may be thought of as the first level of a Lasserre hierarchy). These have been applied succesfully to geometrical problems on S^{n-1} and $\mathbb{R}^n$ (see, e.g., [**BV08**] and [**dLMdOFV21**]), although it is only known how to compute the resulting bounds for small values of k. They may also be applied to bound the size of codes on the binary cube, see, e.g, [**Sch05**], [**Lau07b**], [**Pol19**]. The k-point approach can probably be adapted to yield bounds for the parameters we consider in this chapter as well, especially for $\alpha(\mathbb{B}^n, 3, s)$ (with $k = 3$). We are not aware of the relation between such bounds and the ones we have derived here, but it would be interesting to compare them. Furthermore, it would be interesting to see whether the k-point approach leads to bounds with nice analytical expressions similar to the ones we have seen in this chapter.

Acknowledgments. We would like to thank Christine Bachoc for helpful discussions and comments at an early stage of this work.

Bibliography

[AAR99] G.E. Andrews, R. Askey, and R. Roy. *Special Functions*, volume 71 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1999.

[ABH+05] S. Arora, E. Berger, E. Hazan, G. Kindler, and M. Safra. On non-approximability for quadratic programs. *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 206–215, 2005.

[AL12] M.F. Anjos and J.B. Lasserre, editors. *Handbook on Semidefinite, Conic and Polynomial Optimization*. International Series in Operations Research & Management Science. Springer, Boston, MA, 2012.

[AN04] N. Alon and A. Naor. Approximating the cut-norm via Grothendieck's inequality. *36th annual ACM Symposium on Theory of Computing*, pages 72–80, 2004.

[BBC04] N. Bansal, A. Blum, and S. Chawla. Correlation clustering. *Mach. Learn.*, 46:89–113, 2004.

[BDdOFV14] C. Bachoc, P.E.B. DeCorte, F.M. de Oliveira Filho, and F. Vallentin. Spectral bounds for the independence ratio and the chromatic number of an operator. *Isr. J. Math.*, 202:227–254, 2014.

[BF87] T. Bonnesen and W. Fenchel. *Theory of convex bodies*. BCS Associates, Mountain view, 1987.

[BL06] J. Borwein and A. Lewis. *Convex analysis and nonlinear optimization*. Springer, Berlin, 2006.

[BM21] L. Baldi and B. Mourrain. On moment approximation and the effective Putinar's Positivstellensatz, 2021. `arXiv:2111.11258`.

[BNdOFV09] C. Bachoc, G. Nebe, F.M. de Oliveira Filho, and F. Vallentin. Lower bounds for measurable chromatic numbers. *GAFA*, 19:645–661, 2009.

[BV08] C. Bachoc and F. Vallentin. New upper bounds for kissing numbers from semidefinite programming. *J. Am. Math. Soc.*, 21:909–924, 2008.

[CE03] H. Cohn and N. Elkies. New upper bounds on sphere packings I. *Ann. Math.*, 157:689–714, 2003.

[CLP17] E. Croot, V.F. Lev, and P.P. Pach. Progression-free sets in $\mathbb{Z}_4^n$ are exponentially small. *Ann. Math.*, 185:331–337, 2017.

[CSdOFSV21] D. Castro-Silva, F.M. de Oliveira Filho, L. Slot, and F. Vallentin. A recursive Lovász theta number for simplex-avoiding sets. *To appear in: Proc. Am. Math. Soc.*, 2021. `arxiv:2106.09360`.

[CSdOFSV22] D. Castro-Silva, F.M. de Oliveira Filho, L. Slot, and F. Vallentin. A recursive theta body for hypergraphs, 2022. `arxiv:2206.03929`.

[CV22] S. Christancho and M. Velasco. Harmonic hierarchies for polynomial optimization, 2022. `arxiv:2202.12865`.

[CW04] M. Charikar and A. Wirth. Maximizing quadratic programs: Extending grothendieck's inequality. *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 54–60, 2004.

[dB87] C. de Boor. Cutting corners always works. *Comput. Aided Geom. Des.*, 2:125–131, 1987.

[dBE51] N.G. de Bruijn and P. Erdős. A colour problem for infinite graphs and a problem in the theory of relations. *Indag. Math.*, 13:371–373, 1951.

[dCGHL21] Y. de Castro, F. Gamboa, D. Henrion, and J.B. Lasserre. Dual optimal design and the Christoffel-Darboux polynomial. *Optim. Lett.*, 15(1):3–8, 2021.

[DdOFV20] E. DeCorte, F.M. de Oliveira Filho, and F. Vallentin. Complete positivity and distance-avoiding sets. *Math. Program.*, 2020.

[Del73] P. Delsarte. An algebraic approach to the association schemes of coding theory. *Philips Research Reports Supplements*, 10, 1973.

[Del78] P. Delsarte. Hahn polynomials, discrete harmonics, and t-designs. *SIAM J. Appl. Math.*, 34:157–166, 1978.

[DGS77] P. Delsarte, J.M. Goethals, and J.J. Seidel. Spherical codes and designs. *Geom. Dedicata*, 6:363–388, 1977.

[DH06] L. Dalla and T. Hatziafratis. Strict convexity of sets in analytic terms. *J. Aust. Math. Soc.*, 81:49–61, 2006.

[DJ12] K. Driver and K. Jordaan. Bounds for extreme zeros of some classical orthogonal polynomials. *J. Approx. Theory*, 164(9):1200–1204, 2012.

[dKHL17] E. de Klerk, R. Hess, and M. Laurent. Improved convergence rates for Lasserre-type hierarchies of upper bounds for box-constrained polynomial optimization. *SIAM J. Optim.*, 27(1):347–367, 2017.

[dKL10] E. de Klerk and M. Laurent. Error bounds for some semidefinite programming approaches to polynomial minimization on the hypercube. *SIAM J. Optim.*, 20(6):3104–3120, 2010.

[dKL18] E. de Klerk and M. Laurent. Comparison of Lasserre's measure-based bounds for polynomial optimization to bounds obtained by simulated annealing. *Math. Oper. Res.*, 43(4):1317–1325, 2018.

[dKL19] E. de Klerk and M. Laurent. A survey of semidefinite programming approaches to the generalized problem of moments and their error analysis. In C. Araujo, G. Benkart, C.E. Praeger, and B. Tanbay, editors, *World Women in Mathematics 2018: Proceedings of the First World Meeting for Women in Mathematics (WM)*, pages 17–56, Cham, 2019. Springer International Publishing.

[dKL20a] E. de Klerk and M. Laurent. Convergence analysis of a Lasserre hierarchy of upper bounds for polynomial minimization on the sphere. *Math. Program.*, 2020.

[dKL20b] E. de Klerk and M. Laurent. Worst-case examples for Lasserre's measure–based hierarchy for polynomial optimization on the hypercube. *Math. Oper. Res.*, 45(1):86–98, 2020.

[dKLLS17] E. de Klerk, J.B. Lasserre, M. Laurent, and Z. Sun. Bound-constrained polynomial optimization using only elementary calculations. *Math. Oper. Res.*, 42(3):834–853, 2017.

[dKLS17] E. de Klerk, M. Laurent, and Z. Sun. Convergence analysis for Lasserre's measure-based hierarchy of upper bounds for polynomial optimization. *Math. Program.*, 162:363–392, 2017.

[dKLVS17] E. de Klerk, M. Laurent, J.C. Vera, and Z. Sun. On the convergence rate of grid search for polynomial optimization over the simplex. *Optim. Lett.*, 11:597–608, 2017.

[dLdOFV14] D. de Laat, F.M. de Oliveira Filho, and F. Vallentin. Upper bounds for packings of spheres of several radii. *Forum Math. Sigma*, 2, 2014.

[dLMdOFV21] D. de Laat, F.C. Machado, F.M. de Oliveira Filho, and F. Vallentin. k-Point semidefinite programming bounds for equiangular lines. *Math. Program.*, 2021.

[DN10] D. Dimitrov and G. Nikolov. Sharp bounds for the extreme zeros of classical orthogonal polynomials. *J. Approx. Theory*, 162(10):1793–1804, 2010.

[dOFV10] F.M. de Oliveira Filho and F. Vallentin. Fourier analysis, linear programming, and densities of distance-avoiding sets in $\mathbb{R}^n$. *J. Eur. Math. Soc.*, 12:1417–1428, 2010.

[dOFV19] F.M. de Oliveira Filho and F. Vallentin. Computing upper bounds for the packing density of congruent copies of a convex body. In G. Ambrus, I. Bárány, K.J. Böröczky, G. Fejes Tóth, and J. Pach, editors, *Bolyai Society Mathematical Studies*, volume 27, pages 155–188. Springer, Berlin, Heidelberg, 2019.

[dR47] G. de Rham. Un peu de mathématiques à propos d'une courbe plane. *Elem. Math.*, 2:89–97, 1947.

[DS07] L. Dalla and L. Samiou. Curvature and q-strict convexity. *Beitr. Algebra Geom.*, 48:83–93, 2007.

[Dun78] C.F. Dunkl. An addition theorem for Hahn polynomials: the spherical functions. *SIAM J. Math. Anal.*, 9:627–637, 1978.

[DW13] A.C. Doherty and S. Wehner. Convergence of sdp hierarchies for polynomial optimization on the hypersphere, 2013. `arXiv:1210.5048`.

[DX13] F. Dai and Y. Xu. *Approximation theory and harmonic analysis on spheres and balls*. Springer, Berlin, 2013.

[DX14] C. Dunkl and Y. Xu. *Orthogonal Polynomials of Several Variables*, volume 155 of *Encyclopedia of Mathematics and its Applications*. Cambridge Univ. Press, 2014.

[EG17] J.S. Ellenberg and D. Gijswijt. On large subsets of $\mathbb{F}_q^n$ with no three-term arithmetic progression. *Ann. Math.*, 185:339–343, 2017.

[EL90] A. Elbert and A. Laforgia. Upper bounds for the zeros of ultraspherical polynomials. *J. Approx. Theory*, 61:88–97, 1990.

[Fal81] K.J. Falconer. The realization of distances in measurable subsets covering $\mathbb{R}^n$. *J. Comb. Theory Ser. A*, 31:184 – 189, 1981.

[FF21] K. Fang and H. Fawzi. The sum-of-squares hierarchy on the sphere and applications in quantum information theory. *Math. Program.*, 190:331–360, 2021.

[FGL21] Y. Filmus, K. Golubev, and N. Lifshitz. High dimensional Hoffman bound and applications in extremal combinatorics. *Algebr. Comb.*, 4(6):1005–1026, 2021.

[FR87] P. Frankl and V. Rödl. Forbidden intersections. *Trans. Am. Math. Soc.*, 300:259–286, 1987.

[FSP16] H. Fawzi, J. Saunderson, and P.A. Parrilo. Sparse sums of squares on finite abelian groups and improved semidefinite lifts. *Math. Program.*, 160:149–191, 2016.

[FW81] P. Frankl and R.M. Wilson. Intersection theorems with geometric consequences. *Combinatorica*, 1:357–368, 1981.

[GLS88] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1988.

[Goe97] M. Goemans. Semidefinite programming in combinatorial optimization. *Math. Program.*, 79:143–161, 1997.

[GW95] M. Goemans and D. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42:1115–1145, 1995.

[JH16] C. Josz and D. Henrion. Strong duality in Lasserre's hierarchy for polynomial optimization. *Opt. Lett.*, 10:3–10, 2016.

[Kar72] R.M. Karp. Reducibility among combinatorial problems. In R.E. Miller and J.W. Thatcher, editors, *Complexity of Computer Computations (Proceedings of a symposium on the Complexity of Computer Computations)*, pages 85–103., IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1972. Plenum Press, New York.

[KdK21] F. Kirschner and E. de Klerk. Convergence rates of RLT and Lasserre-type hierarchies for the generalized moment problem over the simplex and the sphere, 2021. `arXiv:2103.02924`.

[KdK22] F. Kirschner and E. de Klerk. Construction of multivariate polynomial approximation kernels via semidefinite programming, 2022. `arXiv:2203.05892`.

[KL12] A. Kroó and D. Lubinsky. Christoffel functions and universality in the bulk for multivariate orthogonal polynomials. *Can. J. Math.*, 65:600–620, 2012.

[KLM16] A. Kurpisz, S. Leppänen, and M. Mastrolilli. Tight sum-of-squares lower bounds for binary polynomial optimization problems. In I. Chatzigiannakis et al., editor, *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 78, pages 1–14, 2016.

[Kou72] D. Koutroufiotis. On Blaschke's rolling theorems. *Arch. Math.*, 23:655–670, 1972.

[Kro15] A. Kroó. Multivariate "needle" polynomials with application to norming sets and cubature formulas. *Acta Math. Hung.*, 147:46–72, 2015.

[KS92] A. Kroó and J. Swetits. On density of interpolation points, a Kadec-type theorem, and Saff's principle of contamination in L_p-approximation. *Constr. Approx.*, 8:87–103, 1992.

[Las01] J.B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2001.

[Las09a] J.B. Lasserre. Approximate volume and integration for basic semialgebraic sets. *SIAM Review*, 51(4):722–743, 2009.

[Las09b] J.B. Lasserre. *Moments, Positive Polynomials and Their Applications*. Imperial College Press, 2009.

[Las11] J.B. Lasserre. A new look at nonnegativity on closed sets and polynomial optimization. *SIAM J. Optim.*, 21(3):864–885, 2011.

[Las20] J.B. Lasserre. Connecting optimization with spectral analysis of tri-diagonal Hankel matrices. *Math. Program.*, 190, 2020.

[Las21] J.B. Lasserre. The moment-SOS hierarchy and the Christoffel-Darboux kernel. *Optim. Lett.*, 15:1835–1845, 2021.

[Lau03a] M. Laurent. A comparison of the Sherali-Adams, Lovász-Schrijver and Lasserre relaxations for 0-1 programming. *Math. Oper. Res.*, 28:470–496, 2003.

[Lau03b] M. Laurent. Lower bound for the number of iterations in semidefinite hierarchies for the cut polytope. *Math. Oper. Res.*, 28:871–883, 2003.

[Lau07a] M. Laurent. Semidefinite representations for finite varieties. *Math. Program*, 109:1–26, 2007.

[Lau07b] M. Laurent. Strengthened semidefinite programming bounds for codes. *Math. Program.*, 109:239–261, 2007.

[Lau09] M. Laurent. Sums of squares, moment matrices and optimization over polynomials. In M. Putinar and S. Sullivant, editors, *Emerging Applications of*

Algebraic Geometry, pages 157–270. Springer New York, New York, NY, 2009.

[Lev98] V.I. Levenshtein. Universal bounds on codes and designs. In *Handbook of Coding Theory*, volume 1, pages 499–648. North-Holland, 1998.

[Lov79] L. Lovász. On the Shannon capacity of a graph. *IEEE Trans. Inf. Theory*, IT-25:1–7, 1979.

[LP19] J.B. Lasserre and E. Pauwels. The empirical Christoffel function with applications in data analysis. *Adv. Comput. Math.*, 45:1439–1468, 2019.

[LS21] M. Laurent and L. Slot. An effective version of Schmüdgen's Positivstellensatz for the hypercube. *Submitted to: Opt. Lett.*, 2021. `arXiv:2109.09528`.

[MPSV20] A. Martinez, F. Piazzon, A. Sommariva, and M. Vianello. Quadrature-based polynomial optimization. *Optim. Lett.*, 14:1027–1036, 2020.

[MPW$^+$21] S. Marx, E. Pauwels, T. Weisser, D. Henrion, and J.B. Lasserre. Semi-algebraic approximation using Christoffel–Darboux kernel. *Constr. Approx.*, 2021.

[MRR78] R.J. McEliece, E.R. Rodemich, and H.C. Rumsey. The Lovász bound and some generalizations. *JCISS*, 3:134–152, 1978.

[MS65] T. S. Motzkin and E. G. Straus. Maxima for graphs and a new proof of a theorem of Turán. *Can. J. Math.*, 17:533 – 540, 1965.

[MS83] F. Macwilliams and N. Sloane. *The Theory of Error Correcting Codes*, volume 16 of *North-Holland Mathematical Library*. Elsevier, 1983.

[Nas20] E. Naslund. Monochromatic equilateral triangles in the unit distance graph. *Bull. London Math. Soc.*, 52:687–692, 2020.

[Nat64] I.P. Natanson. *Constructive Function Theory*, volume 1 of *Uniform Approximation*. University of Michigan Library, 1964.

[NS07] J. Nie and M. Schweighofer. On the complexity of Putinar's Positivstellensatz. *J. Complexity*, 23(1):135–150, 2007.

[Par00] P.A. Parrilo. Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization, 2000. Ph.D. Thesis, California Institute of Technology.

[PD01] A. Prestel and C. Delzell. *Positive polynomials: from Hilbert's 17th problem to real algebra*. Springer, Berlin, Heidelberg, 2001.

[PH13] M.-J. Park and S.-P. Hong. Handelman rank of zero-diagonal quadratic programs over a hypercube and its applications. *J. Glob. Optim.*, 56:727–736, 2013.

[Pol19] S. Polak. New methods in coding theory: Error-correcting codes and the Shannon capacity, 2019. Ph.D. Thesis, University of Amsterdam.

[PPL21] E. Pauwels, M. Putinar, and J.B. Lasserre. Data analysis from empirical moments and the Christoffel function. *Found. Comput. Math.*, 21:243–273, 2021.

[PR00] V. Powers and B. Reznick. Polynomials that are positive on an interval. *Trans. Amer. Math. Soc.*, 352:4677 – 4692, 2000.

[Put93] M. Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana Univ. Math. J.*, 42:969–984, 1993.

[Rez00] B. Reznick. Some concrete aspects of Hilbert's 17th problem. In C. N. Delzell and J.J. Madden, editors, *Real Algebraic Geometry and Ordered Structures*, volume 253 of *Cont. Math.*, pages 251–272. 2000.

[RS75] M. Reed and B. Simon. *Methods of Modern Mathematical Physics II: Fourier Analysis, Self-adjointness*. Academic Press, New York, 1975.

[Sag18a] A.A. Sagdeev. An improved Frankl-Rödl theorem and some of its geometric consequences. *Problemy Peredachi Informatsii*, 54:45–72, 2018.

[Sag18b] A.A. Sagdeev. On a Frankl-Rödl theorem. *Izvestiya Rossiiskoi Akademii Nauk. Seriya Matematicheskaya*, 82:128–157, 2018.

[Sch38] I.J. Schoenberg. Metric spaces and completely monotone functions. *Ann. Math.*, 39:811–841, 1938.

[Sch42] I.J. Schoenberg. Positive definite functions on spheres. *Duke Math. J.*, 9:96–108, 1942.

[Sch79] A. Schrijver. A comparison of the Delsarte and Lovász bounds. *IEEE Trans. Inf. Theory*, IT-25:425–429, 1979.

[Sch91] K. Schmüdgen. The K-moment problem for compact semi-algebraic sets. *Mathematische Annalen*, 289(2):203–206, 1991.

[Sch04] M. Schweighofer. On the complexity of Schmüdgen's Positivstellensatz. *J. Complexity*, 20(4):529–543, 2004.

[Sch05] A. Schrijver. New code upper bounds from the Terwilliger algebra and semidefinite programming. *IEEE Trans. Inf. Theory*, 51(8):2859 – 2866, 2005.

[Sen90] B. Sendov. *Hausdorff approximations*. Kluwer academic publishers, Dordrecht, 1990.

[Sha04] A. Shadrin. Twelve proofs of the Markov inequality. *Approximation theory – a volume dedicated to B. Bojanov*, pages 233–298, 2004.

[SL20] L. Slot and M. Laurent. Improved convergence analysis of Lasserre's measure-based upper bounds for polynomial minimization on compact sets. *Math. Program.*, 193:831–871, 2020.

[SL21a] L. Slot and M. Laurent. Near optimal analysis of Lasserre's univariate measure-based bounds for multivariate polynomial optimization. *Math. Program.*, 188:443–460, 2021.

[SL21b] L. Slot and M. Laurent. Sum-of-squares hierarchies for binary polynomial optimization. In M. Singh and D.P. Williamson, editors, *Integer Programming and Combinatorial Optimization*, pages 43–57, Cham, 2021. Springer International Publishing. *Extended version in Math. Program.*, 2022. `https://doi.org/10.1007/s10107-021-01745-9`.

[Slo22] L. Slot. Sum-of-squares hierarchies for polynomial optimization and the Christoffel-Darboux kernel. *To appear in: SIAM J. Opt.*, 2022. `arxiv:2111.04610`.

[Ste96] G. Stengle. Complexity estimates for the Schmüdgen Positivstellensatz. *J. Complex.*, 12:167–174, 1996.

[STKI17] S. Sakaue, A. Takeda, S. Kim, and N. Ito. Exact semidefinite programming relaxations with truncated moment matrix for binary polynomial optimization problems. *SIAM J. Optim.*, 27(1):565–582, 2017.

[Sze75] G. Szegő. *Orthogonal Polynomials*. American Mathematical Society, 1975.

[Tao16] T. Tao. A symmetric formulation of the Croot-Lev-Pach-Ellenberg-Gijswijt capset bound, 2016. `terrytao.wordpress.com`.

[Tch57] V. Tchakaloff. Formules de cubature mecanique coefficients non négatifs. *Bull. Sci. Math.*, 81:123–134, 1957.

[Val08] F. Vallentin. Lecture notes: Semidefinite programs and harmonic analysis, 2008. `arxiv:0809.2017`.

[VGO+20] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E.W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E.A. Quintero, C.R. Harris, A.M. Archibald, A.H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.

[Wen48] J.G. Wendel. Note on the gamma function. *Amer. Math. Monthly*, 55:563–564, 1948.

[WP86] W. Pleśniak W. Pawłucki. Markov's inequality and C^∞ functions on sets with polynomial cusps. *Math. Ann.*, 275:467–480, 1986.

[WWAF06] A. Weisse, G. Wellein, A. Alvermann, and H. Fehske. The kernel polynomial method. *Rev. Mod. Phys.*, 78:275–306, 2006.

[Xu98] Y. Xu. Summability of Fourier orthogonal series for Jacobi weight functions on the simplex in $\mathbb{R}^d$. *Proc. Amer. Math. Soc.*, 126(10):3027–3036, 1998.

[Xu99] Y. Xu. Summability of Fourier orthogonal series for Jacobi weight on a ball in $\mathbb{R}^d$. *Trans. Amer. Math. Soc.*, 351(6):2439–2458, 1999.

Index

List of symbols

$x, y \in \mathbb{R}$	Variables.
$\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$	n-variate variables.
$\mathbf{x} \cdot \mathbf{y}$, $\langle \mathbf{x}, \mathbf{y} \rangle$	Standard inner product of $\mathbf{x}$ and $\mathbf{y}$.
$\|\mathbf{x}\|$	Euclidean norm of $\mathbf{x}$.
$\mathbb{R}[x], \mathbb{R}[\mathbf{x}]$	Univariate and multivariate polynomial ring.
$\mathbb{R}[x]_d, \mathbb{R}[\mathbf{x}]_d$	Polynomial rings truncated at degree d.
$\mathbb{R}^{n\times n}$	Space of $n \times n$ real matrices.
$\mathcal{S}^n_+$	Space of $n \times n$ positive semidefinite matrices.
$\langle A, B \rangle$	Trace inner product on $\mathbb{R}^{n\times n}$.
$[n]$	Integers $1, 2, \ldots, n$.
$[a : b]$	Integers $a, a+1, \ldots, b$.

Special sets

$\mathbf{X}$	Feasible region for polynomial optimization problems
$\mathbb{B}^n$	The binary hypercube $\{0,1\}^n$.
B^n	The unit ball $\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|^2 \leq 1\}$.
$B^n_\rho(c)$	The ball of radius ρ centered at c.
Δ^n	The standard simplex $\{\mathbf{x} \in \mathbb{R}^n : \mathbf{x} \geq 0,\ \sum_{i=1}^n \mathbf{x}_i \leq 1\}$.
S^{n-1}	The unit sphere $\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|^2 = 1\}$.

Polynomial optimization

f	Objective polynomial to be minimized on $\mathbf{X}$.
$f_{\min}$	Minimum of f on $\mathbf{X}$.
$\mathbf{x}^*$	Minimizer of f on $\mathbf{X}$.
$\|f\|_{\mathbf{X}}$, $\|f\|_\infty$	Supremum norm of f on $\mathbf{X}$ (if no ambiguity is possible).
$\mathcal{P}_+(\mathbf{X})$, $\mathcal{P}_+(\mathbf{X})_r$	Nonnegative polynomials on $\mathbf{X}$ (truncated at degree r).

Sum-of-squares hierarchies

$\Sigma[\mathbf{x}], \Sigma[\mathbf{x}]_r$	Sum-of-squares polynomials (truncated at degree r).
$\mathcal{Q}(\mathbf{X}), \mathcal{Q}(\mathbf{X})_r$	Quadratic module (truncated at degree r).
$\mathcal{T}(\mathbf{X}), \mathcal{T}(\mathbf{X})_r$	Preordering (truncated at degree r).
$\mathcal{M}(\mathbf{X})$	Cone of positive Borel measures supported on $\mathbf{X}$.
$\mathrm{lb}(f, \mathcal{Q}(\mathbf{X}))_r$	Putinar-type lower bound on $f_{\min}$.
$\mathrm{lb}(f, \mathcal{T}(\mathbf{X}))_r$	Schmüdgen-type lower bound on $f_{\min}$.
$\mathrm{ub}(f, \mathbf{X}, \mu)_r$	Measure-based upper bound on $f_{\min}$.
$\mathrm{ub}(f, \mathcal{Q}(\mathbf{X}), \mu)_r$	Measure-based Putinar-type upper bound on $f_{\min}$.
$\mathrm{ub}(f, \mathcal{T}(\mathbf{X}), \mu)_r$	Measure-based Schmüdgen-type upper bound on $f_{\min}$.
$\mathrm{ub}(f, \mathbf{X}, \mu)_r^{\mathrm{pf}}$	Push-forward measure-based upper bound on $f_{\min}$.

Polynomial kernels

$\mathrm{K}(\cdot, \cdot)$	A kernel on $\mathbf{X}$.
$\mathbf{K}(\cdot)$	The linear operator associated to K.
$\mathrm{K}_r^{\mathrm{jac}}(\cdot, \cdot)$	The Jackson kernel of degree r.
$\mathrm{CD}_r(\cdot, \cdot)$	The Christoffel-Darboux kernel of degree r.
$\mathrm{CD}_r(\cdot, \cdot\,; \lambda)$	The perturbed Christoffel-Darboux kernel of degree r.
$\mathrm{lb}(f, \mathbf{X}, \mathrm{K})_{\mathrm{harm}}$	The harmonic lower bound on $f_{\min}$ w.r.t. the kernel K.

Cubature rules

$\mathcal{W} \subseteq \mathbf{X} \times \mathbb{R}$	A (positive) cubature rule for $(\mathbf{X}, \mu)$.
$\mathrm{ub}(f, \mathbf{X}, \mathcal{W})_{\mathrm{cub}}$	The cubature-based upper bound on $f_{\min}$.
$\mathrm{lb}(f, \mathbf{X}, \mathcal{W}, \mathrm{K})_{\mathrm{cub}}$	The cubature-based lower bound on $f_{\min}$ w.r.t. the kernel K.

Orthogonal polynomials

$\mathcal{P}(\mathbf{X})$	Quotient ring of polynomials restricted to $\mathbf{X}$.
$P_\alpha(\cdot)$	Orthonormal polynomial of degree $\lvert\alpha\rvert$ for $\mathcal{P}(\mathbf{X})$ w.r.t a measure μ.
H_k	Subspace of $\mathcal{P}(\mathbf{X})$ spanned by the P_α of degree $\lvert\alpha\rvert = k$.
$P_k(\cdot)$	Orthogonal polynomial of degree k on $I \subseteq \mathbb{R}$ w.r.t. a measure μ.
$\overline{P}_k(\cdot)$	Normalization of P_k satisfying $\max_{x \in I} \lvert P_k(x)\rvert = 1$.
$\widetilde{P}_k(\cdot)$	Normalization of P_k satisfying $\langle P_k, P_k \rangle_\mu = 1$.
$\mathcal{G}_k^{(\alpha)}(\cdot)$	Gegenbauer polynomials.
$\mathcal{J}_k^{(\alpha,\beta)}(\cdot)$	Jacobi polynomials.
$\mathcal{K}_k^{(n)}(\cdot)$	Krawtchouk polynomials.
$\mathcal{C}_k(\cdot)$	Chebyshev polynomials.

Needle polynomials and (convex) geometry

$\nu_r^h(\cdot)$	Needle polynomial of degree r with parameter h.
$\widehat{\nu}_r^h(\cdot)$	$\frac{1}{2}$-needle polynomial of degree r with parameter h.
$\mathrm{int}(\mathbf{X})$	Interior of $\mathbf{X}$.
T_a	Tangent hyperplane at $a \in \mathbf{X}$.
$N_{\mathbf{x}}(\mathbf{X})$	Normal cone of $\mathbf{x} \in \mathbf{X}$.
$v(\mathbf{x}) \in N_{\mathbf{x}}(\mathbf{X})$	Normal vector of $\mathbf{x} \in \mathbf{X}$.
$f \leq_a g$	g is an upper estimator for f exact at $a \in \mathbf{X}$.
$w_1 \succeq w_2, w_1 \succeq_a w_2$	See Section 3.1.5.

Graph theory

$G = (V, E)$	Graph with vertex set V and edge set E.
$\overline{G}$	Complement of G.
$\mathbf{x}_S, \chi_S$	Incidence vector, indicator function of $S \subseteq V$.
$\alpha(G)$	Stability number of G.
$\chi(G)$	Chromatic number of G.
$\vartheta(G)$	Theta number of G.
$\overline{\delta}(X)$	Upper density of $X \subseteq \mathbb{R}^n$.
$\alpha(S^{n-1}, k, t)$	Largest size of a (k, t)-simplex avoiding set in S^{n-1}.
$\alpha(\mathbb{R}^n, k)$	Largest upper density of a k-simplex avoiding set in $\mathbb{R}^n$.
$\vartheta(S^{n-1}, k, t)$	Recursive theta number for $\alpha(S^{n-1}, k, t)$.
$\vartheta(\mathbb{R}^n, k)$	Recursive theta number for $\alpha(\mathbb{R}^n, k)$.

CENTER DISSERTATION SERIES

CentER for Economic Research, Tilburg University, the Netherlands

No.	Author	Title	ISBN	Published
638	Pranav Desai	Essays in Corporate Finance and Innovation	978 90 5668 639 0	January 2021
639	Kristy Jansen	Essays on Institutional Investors, Asset Allocation Decisions, and Asset Prices	978 90 5668 640 6	January 2021
640	Riley Badenbroek	Interior Point Methods and Simulated Annealing for Nonsymmetric Conic Optimization	978 90 5668 641 3	February 2021
641	Stephanie Koornneef	It's about time: Essays on temporal anchoring devices	978 90 5668 642 0	February 2021
642	Vilma Chila	Knowledge Dynamics in Employee Entrepreneurship: Implications for parents and offspring	978 90 5668 643 7	March 2021
643	Minke Remmerswaal	Essays on Financial Incentives in the Dutch Healthcare System	978 90 5668 644 4	July 2021
644	Tse-Min Wang	Voluntary Contributions to Public Goods: A multi-disciplinary examination of prosocial behavior and its antecedents	978 90 5668 645 1	March 2021
645	Manwei Liu	Interdependent individuals: how aggregation, observation, and persuasion affect economic behavior and judgment	978 90 5668 646 8	March 2021
646	Nick Bombaij	Effectiveness of Loyalty Programs	978 90 5668 647 5	April 2021
647	Xiaoyu Wang	Essays in Microeconomics Theory	978 90 5668 648 2	April 2021
648	Thijs Brouwer	Essays on Behavioral Responses to Dishonest and Anti-Social Decision-Making	978 90 5668 649 9	May 2021
649	Yadi Yang	Experiments on hold-up problem and delegation	978 90 5668 650 5	May 2021
650	Tao Han	Imperfect information in firm growth strategy: Three essays on M&A and FDI activities	978 90 5668 651 2	June 2021

No.	Author	Title	ISBN	Published
651	Johan Bonekamp	Studies on labour supply, spending and saving before and after retirement	978 90 5668 652 9	June 2021
652	Hugo van Buggenum	Banks and Financial Markets in Microfounded Models of Money	978 90 5668 653 6	August 2021
653	Arthur Beddock	Asset Pricing with Heterogeneous Agents and Non-normal Return Distributions	978 90 5668 654 3	September 2021
654	Mirron Adriana Boomsma	On the transition to a sustainable economy: Field experimental evidence on behavioral interventions	978 90 5668 655 0	September 2021
655	Roweno Heijmans	On Environmental Externalities and Global Games	978 90 5668 656 7	August 2021
656	Lenka Fiala	Essays in the economics of education	978 90 5668 657 4	September 2021
657	Yuexin Li	Pricing Art: Returns, Trust, and Crises	978 90 5668 658 1	September 2021
658	Ernst Roos	Robust Approaches for Optimization Problems with Convex Uncertainty	978 90 5668 659 8	September 2021
659	Joren Koëter	Essays on asset pricing, investor preferences and derivative markets	978 90 5668 660 4	September 2021
660	Ricardo Barahona	Investor Behavior and Financial Markets	978 90 5668 661 1	October 2021
660	Stefan ten Eikelder	Biologically-based radiation therapy planning and adjustable robust optimization	978 90 5668 662 8	October 2021
661	Maciej Husiatyński	Three essays on Individual Behavior and New Technologies	978 90 5668 663 5	October 2021
662	Hasan Apakan	Essays on Two-Dimensional Signaling Games	978 90 5668 664 2	October 2021
663	Ana Moura	Essays in Health Economics	978 90 5668 665 9	November 2021
664	Frederik Verplancke	Essays on Corporate Finance: Insights on Aspects of the General Business Environment	978 90 5668 666 6	October 2021

No.	Author	Title	ISBN	Published
665	Zhaneta Tancheva	Essays on Macro-Finance and Market Anomalies	978 90 5668 667 3	November 2021
666	Claudio Baccianti	Essays in Economic Growth and Climate Policy	978 90 5668 668 0	November 2021
667	Hongwei Zhang	Empirical Asset Pricing and Ensemble Machine Learning	978 90 5668 669 7	November 2021
668	Bart van der Burgt	Splitsing in de Wet op de vennootschapsbelasting 1969 Een evaluatie van de Nederlandse winstbelastingregels voor splitsingen ten aanzien van lichamen	978 90 5668 670 3	December 2021
669	Martin Kapons	Essays on Capital Markets Research in Accounting	978 90 5668 671 0	December 2021
670	Xolani Nghona	From one dominant growth mode to another: Switching between strategic expansion modes	978 90 5668 672 7	December 2021
671	Yang Ding	Antecedents and Implications of Legacy Divestitures	978 90 5668 673 4	December 2021
672	Joobin Ordoobody	The Interplay of Structural and Individual Characteristics	978 90 5668 674 1	February 2022
673	Lucas Avezum	Essays on Bank Regulation and Supervision	978 90 5668 675 8	March 2022
674	Oliver Wichert	Unit-Root Tests in High-Dimensional Panels	978 90 5668 676 5	April 2022
675	Martijn de Vries	Theoretical Asset Pricing under Behavioral Decision Making	978 90 5668 677 2	June 2022
676	Hanan Ahmed	Extreme Value Statistics using Related Variables	978 90 5668 678 9	June 2022
677	Jan Paulick	Financial Market Information Infrastructures: Essays on Liquidity, Participant Behavior, and Information Extraction	978 90 5668 679 6	June 2022
678	Freek van Gils	Essays on Social Media and Democracy	978 90 5668 680 2	June 2022

No.	Author	Title	ISBN	Published
679	Suzanne Bies	Examining the Effectiveness of Activation Techniques on Consumer Behavior in Temporary Loyalty Programs	978 90 5668 681 9	July 2022
680	Qinnan Ruan	Management Control Systems and Ethical Decision Making	978 90 5668 682 6	June 2022
681	Lingbo Shen	Essays on Behavioral Finance and Corporate Finance	978 90 5668 683 3	August 2022
682	Joshua Eckblad	Mind the Gales: An Attention-Based View of Startup Investment Arms	978 90 5668 684 0	August 2022
683	Rafael Greminger	Essays on Consumer Search	978 90 5668 685 7	August 2022
684	Suraj Upadhyay	Essay on policies to curb rising healthcare expenditures	978 90 5668 686 4	September 2022
685	Bert-Jan Butijn	From Legal Contracts to Smart Contracts and Back Again: An Automated Approach	978 90 5668 687 1	September 2022
686	Sytse Duiverman	Four essays on the quality of auditing: Causes and consequences	978 90 5668 688 8	October 2022
687	Lucas Slot	Asymptotic Analysis of Semidefinite Bounds for Polynomial Optimization and Independent Sets in Geometric Hypergraphs	978 90 5668 689 5	September 2022

The goal of a mathematical optimization problem is to maximize an objective (or minimize a cost) under a given set of rules, called constraints. Optimization has many applications, both in other areas of mathematics and in the real world. Unfortunately, some of the most interesting problems are also very hard to solve numerically. To work around this issue, one often considers relaxations: approximations of the original problem that are much easier to solve. Naturally, it is then important to understand how (in)accurate these relaxations are.

This thesis consists of three parts, each covering a different method that uses semidefinite programming to approximate hard optimization problems.
In Part 1 and Part 2, we consider two hierarchies of relaxations for polynomial optimization problems based on sums of squares. We show improved guarantees on the quality of Lasserre's measure-based hierarchy in a wide variety of settings (Part 1). We establish error bounds for the moment-SOS hierarchy in certain fundamental special cases. These bounds are much stronger than the ones obtained from existing, general results (Part 2).
In Part 3, we generalize the celebrated Lovász theta number to (geometric) hypergraphs. We apply our generalization to formulate relaxations for a type of independent set problem in the hypersphere. These relaxations allow us to improve some results in Euclidean Ramsey theory.

Lucas Slot (Amsterdam, The Netherlands, 1996) received his bachelor's degrees in Mathematics (with honours) and Computer Science from the University of Amsterdam in 2016. He obtained his master's degree in Mathematics from the University of Bonn in 2018.

ISBN: 978 905668 689 5
DOI: 10.26116/wsbs-kt84

TILBURG SCHOOL OF ECONOMICS
AND MANAGEMENT

CentER Dissertation Series