

DOCTORAL THESIS

Parameter recovery for moment problems on algebraic varieties

Markus Wageringel

March 2022

Universität Osnabrück
Fachbereich Mathematik/Informatik
Advisor: Stefan Kunis

Contents

Introduction	4
1 Preliminaries	8
1.1 Terminology	8
1.2 Algebraic prerequisites	9
1.3 Prony’s method	11
1.4 Special case of equal weights	17
1.5 Connection to symmetric tensor decomposition	18
2 Local Dirac mixtures	26
2.1 Moment variety	27
2.2 Pareto distribution	28
2.3 Recovery of parameters	30
2.4 Applications	37
3 Moment problems on positive-dimensional varieties	43
3.1 Sesquilinearity	44
3.2 Factorization properties	50
3.3 Considerations for the complex torus	54
3.4 Recovery of the support from moments	56
3.5 From moments to approximations of measures	63
4 Recovery of components from eigenvalues	79
4.1 Prologue	80
4.2 Generalized eigenvalues	82
4.3 Hermitian and symmetric cases	93
4.4 Positive-semidefinite case	107
4.5 A sufficient condition	112
4.6 Sharpness results	118
4.7 Revised algorithms	122
4.8 Numerical implementation	130
List of figures	142
List of algorithms	142
Glossary of symbols	143
Bibliography	145

Introduction

One of the main topics in signal processing is the reconstruction of a signal from a set of measurements. Similarly, in statistics, much interest pertains to finding a ground truth explaining a sample set. A possible set of measurements are the *moments* of a measure. These are quantities that reflect aspects such as the shape of a distribution. For a Borel measure μ (or, more generally, a distribution) on the space \mathbb{R}^n , $n \in \mathbb{N}$, its moments are defined as

$$m_\alpha := \int_{\mathbb{R}^n} x^\alpha d\mu(x), \quad \alpha \in \mathbb{N}^n,$$

where $x^\alpha := x_1^{\alpha_1} \cdots x_n^{\alpha_n}$. Moments also arise in signal processing as the Fourier coefficients of a signal.

Of particular interest for this thesis is the *truncated moment problem*: Is it possible to reconstruct the measure μ from a truncated moment sequence, i. e. from finitely many moments $\{m_\alpha \mid |\alpha| \leq d\}$, for some $d \in \mathbb{N}$, and how can this be accomplished effectively? For this, we restrict ourselves to particular parametric classes of measures and seek to solve the inverse problem of recovering the defining parameters of the underlying measure.

One such class is that of *finitely-supported* (complex-signed) measures of the form

$$\sum_{j=1}^r \lambda_j \delta_{\xi_j},$$

supported at points $\xi_1, \dots, \xi_r \in \mathbb{C}^n$ with weights $\lambda_1, \dots, \lambda_r \in \mathbb{C} \setminus \{0\}$, for some $r \in \mathbb{N}$, where δ_{ξ_j} denotes the Dirac measure located at the point ξ_j . For a fixed number of support points r , the parameters $(\xi_1, \lambda_1), \dots, (\xi_r, \lambda_r)$ are in fact uniquely determined by the moments $\{m_\alpha \mid |\alpha| \leq d\}$, as long as $d \in \mathbb{N}$ is sufficiently large.

In the one-dimensional case $n = 1$, this parameter recovery problem can be solved using *Prony's method*, a widely-used tool in signal processing that is algebraic at heart. It is precisely this algebraic nature that we focus on and use as motivation for generalizations. The classic Prony method and its modern multivariate form recover the finitely many support points as the zero set of a family of polynomials. Therefore, it is natural to view the support as a zero-dimensional algebraic variety. In the course of the thesis, we not only consider finitely-supported measures or distributions, but we also examine measures which have finite moments and a support that is (contained in) a positive-dimensional variety, such as an algebraic curve or surface. As these varieties consist of infinitely many

points, this forms a large class of measures and we study related reconstruction problems for them.

Outline

In Chapter 1, we start by giving an overview of the terminology we use and by summarizing a few algebraic and ideal theoretic properties. Afterwards, we give an introduction to the multivariate Prony method which can be used to solve the parameter recovery problem for finitely-supported signed measures in any dimension. It serves as common starting point for the following chapters of the thesis. Additionally, a multi-degree version of Prony's method is presented which, in particular, is useful for the demonstration of a link to the decomposition of symmetric tensors, at the end of the chapter.

Chapter 2 is an adaptation of [GW20] and deals with local Dirac mixture distributions. These can be regarded as a generalization of finitely-supported measures that incorporate information about derivatives. We summarize results about the moment variety of these distributions and relate it to the moment variety of Pareto distributions. Beyond that, we explore the parameter recovery problem. By connecting it to Prony's method, we formulate a numerical reconstruction algorithm, which is then applied to some examples.

In Chapter 3, we switch from finitely-supported measures to the much more general class of measures that are supported on algebraic varieties of any dimension. We analyze which features of Prony's method can be transferred to this setting. By considering the kernels of certain moment matrices, it is possible to recover the vanishing ideal of the support of a measure. In other words, given sufficiently many moments, one obtains the Zariski closure of the support, by algebraic means. Afterwards, we focus on measures whose support is contained in the complex torus. Besides the identification of the support, we examine functions that are constructed from finitely many moments and that approximate particular features of the underlying measure on the torus.

Finally, Chapter 4 addresses the problem of finding the individual components of mixtures of measures that are supported on different algebraic varieties. As in the case of finitely-supported measures, this problem can be approached by eigenvalue-based methods. For measures on positive-dimensional varieties, this strategy requires some alterations which we discuss in detail. We present several reconstruction algorithms that allow to solve this problem. An implementation of the algorithms, both symbolic and numerical, is provided at [Wag21].

Contributions

The main contributions of this thesis are the following:

- Algorithm 2.1 represents a procedure for reconstructing l -th order local Dirac mixtures, which is proved in Theorem 2.3.2. Our setting is a special case of the one considered by [Mou18], which allows us to optimize the algorithm for the use of

fewer moments than would be needed in a more general context.

- The Vandermonde decomposition of the moment matrix of a finitely-supported signed measure is an essential ingredient of Prony’s method. Theorem 3.2.4 forms an analog of this decomposition that is suitable also for measures supported on positive-dimensional varieties.
- For any compactly-supported signed measure, Theorem 3.4.3 establishes a relationship between moment matrices and the vanishing ideal of the support. It shows that the Zariski closure of the support can be computed from finitely many moments. The theorem can be viewed as an extension of Theorem 3.4.11, which makes a similar statement for non-negative measures and has been considered by [LR12; PPL21], in the real affine (non-trigonometric) setting.
- Theorem 3.5.13 proves a pointwise convergence property of certain functions associated to the moments of measures supported on algebraic varieties of any dimension. In the finitely-supported case, this is connected to MUSIC [Sch86], a well-known parameter estimation technique, so the theorem is a generalization for particular measures with infinite support.
- The development of Algorithms 4.2 to 4.4 forms the core of Chapter 4. Their correctness is proved in Theorems 4.4.7, 4.7.4 and 4.7.9, respectively. Algorithm 4.2 solves a parameter recovery problem for pencils of positive-semidefinite matrices. Algorithms 4.3 and 4.4, on the other hand, address the task of finding the components of mixtures of measures supported on positive-dimensional varieties, putting more emphasis on the algebraic geometric nature of the problem – for this, Theorem 4.5.4 is of central importance, by providing a characterization of specific eigenvectors.

All these algorithms are eigenvalue-based. As such, they represent a variant of pencil-based techniques, such as ESPRIT, that are in common use for the case of finitely-supported signed measures [RK90; HS90; ACdH10; Moi15]. By allowing non-discrete measures, our setting is also a more general form of the one considered in multi-snapshot spectral estimation, which appears in problems like direction-of-arrival estimation (see e. g. [KV96; LZGL21]).

Acknowledgments

First and foremost, I thank my supervisor Stefan Kunis and cosupervisor Tim Römer for their guidance and support, over the years. Additionally, I would like to thank Timo de Wolff for reviewing the thesis. My research at Osnabrück was partially supported by the DFG grant GRK 1916 for which I am grateful.

Moreover, I would like to give thanks to my wonderful coauthors Arthur Bik, Adam Czapliński and Alex Grosdos. I thank Paul Breiding for instructions on using the homotopy continuation Julia package as well as the SageMath developers for their continued support and development of the software which was invaluable for my research.

I would like to thank my friends and colleagues for their support and continued friendship during the last years, which include in particular Alex, Arun, Carina, Dominik, Jan, Jonathan, Lorenzo, Mathias, Paul, Robert, Stephan, Tarek and Ulrich. Finally, I thank my family for the enduring support and being there every step of the way.

1 Preliminaries

1.1 Terminology

Let us fix some notation and terminology that is used throughout this thesis. The symbol \mathbb{k} always denotes a field. Note that in some sections we explicitly assume that the field \mathbb{k} satisfies additional properties, such as being of characteristic 0. The natural numbers are denoted by $\mathbb{N} = \{0, 1, 2, \dots\}$.

We choose the following conventions from linear algebra. The standard basis of the vector space \mathbb{k}^n , $n \in \mathbb{N}$, is denoted by e_1, \dots, e_n . If $v \in \mathbb{k}^n$ is a vector, then the coordinates are denoted by v_1, \dots, v_n , with respect to the standard basis or some other specified basis. The homogeneous coordinates of v are written as $[v] = [v_1 : \dots : v_n] \in \mathbb{P}_{\mathbb{k}}^{n-1}$. If no confusion is possible, we may identify a matrix with its corresponding linear map, using the convention that vectors are multiplied to matrices from the right. Thus, the kernel of a matrix is the same as the right null space. Similar conventions apply to the image and related notions. If $V \subseteq \mathbb{k}^n$ is a vector subspace and no confusion is possible, we may also denote by V a matrix whose columns span this space. The trivial vector space is denoted by 0. The (algebraic) dual space of a \mathbb{k} -vector space V is written as V^* or $\text{Hom}_{\mathbb{k}}(V, \mathbb{k})$. Similarly, $\text{Hom}_{\mathbb{k}}^{\text{semi}}(V, \mathbb{k})$ denotes the set of semilinear maps from V to \mathbb{k} (cf. Section 3.1). The notion of positive-semidefiniteness of a matrix only applies to matrices that are also Hermitian.

For an introduction to algebraic geometry, we refer to [CLO15]. For more advanced topics, see also [Har87] and [Eis99]. In this thesis, the term *algebraic variety* refers to the vanishing set of a set of polynomials, also known as algebraic set, that is, we do not require irreducibility. A variety generated by an ideal \mathfrak{a} is denoted by $V(\mathfrak{a})$. The vanishing ideal of a set $X \subseteq \mathbb{k}^n$ is denoted by $I(X)$. Over an infinite field, a property holds *generically*, if it is satisfied for all elements of a non-empty Zariski-open set.

Additionally, let us fix some algebraic notions. Unless stated otherwise, all rings are unital. We use multi-index notation for monomials. Thus, when working in the polynomial ring $\mathbb{k}[x_1, \dots, x_n]$, the monomials are denoted by $x^\alpha = x_1^{\alpha_1} \cdots x_n^{\alpha_n}$, $\alpha \in \mathbb{N}^n$. The *(total) degree* of a polynomial $p = \sum_{\alpha \in \mathbb{N}^n} p_\alpha x^\alpha$ with coefficients $p_\alpha \in \mathbb{k}$ is given by $\deg(p) = \max\{|\alpha| \mid \alpha \in \mathbb{N}^n, p_\alpha \neq 0\}$, where $|\alpha| := \alpha_1 + \dots + \alpha_n$. Similarly, we define the *max-degree* of a Laurent polynomial $q = \sum_{\alpha \in \mathbb{Z}^n} q_\alpha x^\alpha$, $q_\alpha \in \mathbb{k}$, as $\max\{|\alpha|_\infty \mid \alpha \in \mathbb{Z}^n, q_\alpha \neq 0\}$, where $|\alpha|_\infty := \max\{|\alpha_1|, \dots, |\alpha_n|\}$. The same definition applies when q is a polynomial. Though, note that the max-degree does not define a grading of the polynomial ring, but gives rise to a filtration (cf. Example 3.1.3). By $\langle - \rangle$, we denote

the vector subspace spanned by a family of vectors or the ideal spanned by a family of ring elements. If $B \subseteq \mathbb{k}[x_1, \dots, x_n]$ is a family of linearly independent polynomials, we also write $\bigoplus_{f \in B} \mathbb{k}f$ for the vector subspace generated by these elements, in order to emphasize that B is chosen as basis and to avoid ambiguity.

Given an ideal $\mathfrak{a} \subseteq \mathbb{k}[x_1, \dots, x_n]$, the Krull-dimension of the quotient ring $\mathbb{k}[x_1, \dots, x_n]/\mathfrak{a}$, i. e. the supremum of the heights of all prime ideals, is the same as the dimension of the variety $V(\mathfrak{a}) \subseteq \mathbb{k}^n$ (cf. [CLO15, Theorem 9.3.8]). By abuse of language, we also refer to this as the dimension of the ideal \mathfrak{a} . The residue class of a polynomial $p \in \mathbb{k}[x_1, \dots, x_n]$ modulo an ideal \mathfrak{a} is denoted by $\bar{p} = p + \mathfrak{a}$ or, if there is no risk of confusion, it may also be denoted by p again. (Note also that complex conjugation of a vector $v \in \mathbb{C}^n$ is denoted by \bar{v} .) We write

$$\mathfrak{m}_\xi := \langle x - \xi \rangle = \langle x_1 - \xi_1, \dots, x_n - \xi_n \rangle$$

for the maximal ideal associated to a point $\xi \in \mathbb{k}^n$. Furthermore, the map

$$\text{ev}_\xi: \mathbb{k}[x_1, \dots, x_n] \longrightarrow \mathbb{k}, \quad p \longmapsto p(\xi),$$

denotes the *evaluation homomorphism* associated to the point $\xi \in \mathbb{k}^n$. It can naturally be viewed as a ring homomorphism to the quotient ring corresponding to the ideal \mathfrak{m}_ξ .

Unless otherwise noted, the term *measure* refers to non-negative Borel measures. Occasionally, we also work with signed measures, which will especially be common for the finitely-supported measures we consider. Over the complex numbers, the term *signed measure* stands for complex(-signed) measure. Every (finite) non-negative measure is, in particular, a signed measure. Moreover, signed measures are particular *distributions*. Distributions will play a role primarily in Chapter 2. For details, we refer to [Sch73; Rud87].

1.2 Algebraic prerequisites

For completeness, we collect a few elementary ideal theoretic and geometric properties which will be useful for our work later on. We start with some properties of comaximal ideals. Recall that two ideals $\mathfrak{a}, \mathfrak{b}$ of a ring R are *comaximal* if $\mathfrak{a} + \mathfrak{b} = R$.

Lemma 1.2.1. *Let $\mathfrak{a}, \mathfrak{b}$ be comaximal ideals of a ring R . Assume that $\mathfrak{a}'\mathfrak{b} \subseteq \mathfrak{a}$ for some ideal $\mathfrak{a}' \subseteq R$. Then $\mathfrak{a}' \subseteq \mathfrak{a}$.*

Proof. As \mathfrak{a} and \mathfrak{b} are comaximal, we can choose elements $a \in \mathfrak{a}, b \in \mathfrak{b}$ such that $a + b = 1$. Then, for every element $f \in \mathfrak{a}'$, we have $f = f(a + b) = fa + fb \in \mathfrak{a}$, since $fa \in \mathfrak{a}$ and $fb \in \mathfrak{a}'\mathfrak{b} \subseteq \mathfrak{a}$. \square

The following statement is a general form of the *Chinese Remainder Theorem*.

Lemma 1.2.2 (cf. [Bou06, Chapter 2.1.2, Proposition 5]). *Let $r \geq 1$ and let $\mathfrak{a}_1, \dots, \mathfrak{a}_r \subseteq R$ be ideals in a ring R which are pairwise comaximal. Then*

- (1) $\mathfrak{a}_1 \cap \cdots \cap \mathfrak{a}_r = \mathfrak{a}_1 \cdots \mathfrak{a}_r$,
- (2) $R/(\mathfrak{a}_1 \cdots \mathfrak{a}_r) \cong R/\mathfrak{a}_1 \oplus \cdots \oplus R/\mathfrak{a}_r$.

Lemma 1.2.3. *Let $r \geq 1$, let $\mathfrak{a}_1, \dots, \mathfrak{a}_r \subseteq R$ be ideals and let $\mathfrak{p} \subseteq R$ be a prime ideal in a ring R such that $\mathfrak{a}_j \not\subseteq \mathfrak{p}$ for $1 \leq j \leq r$. Then*

$$\mathfrak{a}_1 \cdots \mathfrak{a}_r \not\subseteq \mathfrak{p}.$$

In particular, this implies that

$$\bigcap_{j=1}^r \mathfrak{a}_j \setminus \mathfrak{p} \neq \emptyset.$$

Proof. Assume that $\mathfrak{a}_1 \cdots \mathfrak{a}_r \subseteq \mathfrak{p}$. For every $1 \leq j \leq r$, we can choose an element $f_j \in \mathfrak{a}_j \setminus \mathfrak{p}$, since $\mathfrak{a}_j \not\subseteq \mathfrak{p}$. This means that $\prod_{j=1}^r f_j \in \mathfrak{a}_1 \cdots \mathfrak{a}_r \subseteq \mathfrak{p}$. As \mathfrak{p} is a prime ideal, this implies that $f_j \in \mathfrak{p}$ for some j , which is a contradiction by choice of f_j . Therefore, we have $\mathfrak{a}_1 \cdots \mathfrak{a}_r \not\subseteq \mathfrak{p}$. The addendum follows from $\mathfrak{a}_1 \cdots \mathfrak{a}_r \subseteq \mathfrak{a}_1 \cap \cdots \cap \mathfrak{a}_r$. \square

Moreover, let us take note of the following property related to (affine) Hilbert functions (cf. [CLO15, Chapter 9.3]). Here, $\mathcal{O}(f(d))$, $\Theta(f(d))$ denote the sets of functions bounded from above or, respectively, from above and below (up to constant factors) by a function $f(d)$, for all sufficiently large d .

Lemma 1.2.4. *Let $\mathcal{X}_1, \dots, \mathcal{X}_r \subseteq \mathbb{k}^n$ be varieties of dimension $s \in \mathbb{N}$ over a field \mathbb{k} of characteristic 0 such that $\dim(\mathcal{X}_i \cap \mathcal{X}_j) < s$ for all $1 \leq i, j \leq r$ with $i \neq j$. Denote by $\mathfrak{a}_j \subseteq R := \mathbb{k}[x_1, \dots, x_n]$ the vanishing ideal of \mathcal{X}_j and set $\mathfrak{a} := \bigcap_{j=1}^r \mathfrak{a}_j$ as well as $\mathfrak{b}_j := \bigcap_{k=1, k \neq j}^r \mathfrak{a}_k$, for $1 \leq j \leq r$. Then*

- (1) $\dim((\mathfrak{b}_j \cap R_{\leq d})/(\mathfrak{a} \cap R_{\leq d})) \in \Theta(d^s)$, for every $1 \leq j \leq r$;
- (2) $\dim(R_{\leq d}/(\mathfrak{a} \cap R_{\leq d})) - \sum_{j=1}^r \dim((\mathfrak{b}_j \cap R_{\leq d})/(\mathfrak{a} \cap R_{\leq d})) \in \mathcal{O}(d^{s-1})$.

Here, $R_{\leq d}$ denotes the space of polynomials of total degree at most $d \in \mathbb{N}$.

Proof. These statements follow from an analysis of Hilbert polynomials. Indeed, for all $1 \leq j \leq r$, we have

$$\begin{aligned} \dim((\mathfrak{b}_j \cap R_{\leq d})/(\mathfrak{a} \cap R_{\leq d})) &= \dim(\mathfrak{b}_j \cap R_{\leq d}) - \dim(\mathfrak{a} \cap R_{\leq d}) \\ &= -\dim(R_{\leq d}/(\mathfrak{b}_j \cap R_{\leq d})) + \dim(R_{\leq d}/(\mathfrak{a} \cap R_{\leq d})). \end{aligned}$$

The affine Hilbert polynomials of R/\mathfrak{b}_j and R/\mathfrak{a} have leading terms $\frac{d^s}{s!} \sum_{k=1, k \neq j}^r \deg \overline{\mathcal{X}_k}$ and $\frac{d^s}{s!} \sum_{k=1}^r \deg \overline{\mathcal{X}_k}$, respectively, by [Har87, Proposition 1.7.6] combined with [CLO15, Theorem 9.3.12 (ii)], where $\overline{\mathcal{X}_k}$ denotes the closure of \mathcal{X}_k in $\mathbb{P}_{\mathbb{k}}^n$. Thus, the affine Hilbert polynomial of $\mathfrak{b}_j/\mathfrak{a}$ has the leading term $\frac{d^s}{s!} \deg \overline{\mathcal{X}_j}$, from which the statements follow. \square

In particular, if the dimension of the varieties is $s = 0$, then (2) is zero for sufficiently large $d \in \mathbb{N}$, which agrees with the Chinese Remainder Theorem, Lemma 1.2.2.

Furthermore, recall the following property about prime ideals in Laurent polynomial rings.

Lemma 1.2.5. *Let $\mathfrak{a} \subseteq \mathbb{k}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ be an ideal, where $n \in \mathbb{N}$ and \mathbb{k} is any field. Then \mathfrak{a} is prime if and only if $\mathfrak{a} \cap \mathbb{k}[x_1, \dots, x_n]$ is prime.*

Proof. This follows from the fact that the Laurent ring is the localization of the polynomial ring at the multiplicatively closed set $\{x^\alpha\}_{\alpha \in \mathbb{N}^n}$; see for instance [Eis99, Proposition 2.2]. \square

1.3 Prony's method

1.3.1 Multivariate Prony method

The following is a multivariate generalization of *Prony's method* that, in its univariate form, goes back to [Pro95]. It is the central starting point for many of our considerations in this thesis. The variant we cite here is useful for this, but there are many alternative formulations that accentuate different points of view. For instance, it has been considered in terms of exponential sums with a focus on signal processing in [KPRvdO16; vdOhe17; Sau17; Mou18]. Another variation of Prony's method is *Sylvester's algorithm* [Sy186]. It is also related to *Macaulay inverse systems* (see e. g. [Eis99, Chapter 21.2]) and *apolarity theory* (cf. [IK99, Lemma 1.15, algorithm in Chapter 5.4], [Sch17, Chapter 19]), which put more emphasis on algebraic and geometric aspects. We will see further variants later on, in Theorems 1.3.6 and 2.3.1.

Theorem 1.3.1 ([Pro95], [KPRvdO16], [vdOhe17, Corollary 2.19]). *Let \mathbb{k} be a field and let $R = \mathbb{k}[x_1, \dots, x_n]$ be the polynomial ring in n variables. Let $\sigma: R \rightarrow \mathbb{k}$ be a \mathbb{k} -linear map of the form*

$$\sigma(p) = \sum_{j=1}^r \lambda_j p(\xi_j), \quad (1.1)$$

where $\xi_j \in \mathbb{k}^n$ are distinct points and $\lambda_j \in \mathbb{k} \setminus \{0\}$, for $1 \leq j \leq r$. If $d \in \mathbb{N}$ such that $\text{ev}_{\leq d-1}: R_{\leq d-1} \rightarrow \mathbb{k}^r$, $p \mapsto (p(\xi_j))_{j=1}^r$, is surjective, then

$$\mathbb{V}(\ker H_d) = \{\xi_1, \dots, \xi_r\},$$

where $H_d := (\sigma(x^{\alpha+\beta}))_{\alpha, \beta \in \mathbb{N}^n, |\alpha|, |\beta| \leq d}$.

Here, $\ker H_d$ is viewed as a subset of $R_{\leq d}$, the polynomials of degree at most d , by noting that, with respect to the monomial basis, the matrix H_d represents the \mathbb{k} -linear map into the dual space of the vector space $R_{\leq d}$ given by

$$R_{\leq d} \longrightarrow R_{\leq d}^*, \quad p \longmapsto (q \mapsto \sigma(pq)),$$

as well as the \mathbb{k} -bilinear mapping

$$R_{\leq d} \times R_{\leq d} \longrightarrow \mathbb{k}, \quad (p, q) \longmapsto \sigma(pq).$$

Then, $V(\ker H_d)$ denotes the common zero set of all the polynomials that are annihilated by H_d . The ideal spanned by $\ker H_d$ in R , that is, the vanishing ideal of the points $\{\xi_1, \dots, \xi_r\}$, is also called *Prony ideal*.

A map of the form (1.1) can also be viewed as *exponential sum*. It satisfies $\sigma(x^\alpha) = \sum_{j=1}^r \lambda_j \xi_j^\alpha$ for all $\alpha \in \mathbb{N}^n$, so can be interpreted as a map $\mathbb{N}^n \rightarrow \mathbb{k}$, by composing it with $\alpha \mapsto x^\alpha$. Moreover, the map σ is a linear combination of evaluation maps corresponding to the points ξ_1, \dots, ξ_r , so $\sigma = \sum_{j=1}^r \lambda_j \text{ev}_{\xi_j}$, where ev_{ξ_j} denotes the evaluation homomorphism associated to the point ξ_j , i. e. the element of the dual space of R satisfying $\text{ev}_{\xi_j}(p) = p(\xi_j)$ for all $p \in R$, for $1 \leq j \leq r$.

Also note that σ is the *moment functional* of the finitely-supported measure $\mu := \sum_{j=1}^r \lambda_j \delta_{\xi_j}$. For this interpretation, we usually assume that \mathbb{k} is \mathbb{R} or \mathbb{C} . If $\mathbb{k} = \mathbb{C}$ and the weights $\lambda_1, \dots, \lambda_r \in \mathbb{C}$ are complex, then μ is a *signed* (complex) measure, which is explicitly allowed in this setting. Here, δ_{ξ_j} denotes the Dirac measure supported at the point $\xi_j \in \mathbb{k}^n$ for $1 \leq j \leq r$, defined as

$$\delta_{\xi_j}(A) = \begin{cases} 1, & \text{if } \xi_j \in A, \\ 0, & \text{if } \xi_j \notin A, \end{cases}$$

for a set $A \subseteq \mathbb{k}^n$. Indeed, the signed measure μ satisfies $\int_{\mathbb{k}^n} x^\alpha d\mu(x) = \sum_{j=1}^r \lambda_j \xi_j^\alpha = \sigma(x^\alpha)$, so $\sigma(x^\alpha)$ agrees with the α -th moment of μ . On top of that, the moments $\sigma(x^\alpha)$ uniquely determine the map σ . From this point of view, the statement of Theorem 1.3.1 is that the support of the finitely-supported signed measure μ is already determined by *finitely many* of its moments, namely the ones that are required to construct the matrix H_d . In fact, in this case, the weights $\lambda_1, \dots, \lambda_r$ can be recovered as well, by subsequently solving a linear system of equations (cf. [vdOhe17, Algorithm 2.1]), so the measure μ is fully determined by these moments. Note that $\text{ev}_{\leq d-1}$ is surjective if d is sufficiently large, a trivial bound being $d \geq r$; cf. [vdOhe17, Corollary 2.20]. The ideal

$$\bigcap_{j=1}^r \langle x - \xi_j \rangle = \prod_{j=1}^r \langle x - \xi_j \rangle = \prod_{j=1}^r \langle x_1 - \xi_{j1}, \dots, x_n - \xi_{jn} \rangle$$

is clearly generated by polynomials of degree at most r , but in the multivariate setting with $n \geq 2$, unless the points ξ_1, \dots, ξ_r are contained in a one-dimensional subspace of \mathbb{k}^n , this bound can be much larger than necessary.

A more practical sufficient criterion for the evaluation map $\text{ev}_{\leq d-1}$ being surjective is obtained by checking the rank of the matrix H_{d-1} . We can factor H_{d-1} as

$$H_{d-1} = \left(\sum_{j=1}^r \lambda_j \xi_j^{\alpha+\beta} \right)_{|\alpha|, |\beta| \leq d-1} = V^\top \text{diag}(\lambda_1, \dots, \lambda_r) V,$$

where $V = (\xi_j^\alpha)_{1 \leq j \leq r, |\alpha| \leq d-1}$ denotes the *Vandermonde matrix* associated to the points ξ_1, \dots, ξ_r up to degree $d-1$. This matrix corresponds to the linear map

$$\text{ev}_{\leq d-1}: R_{\leq d-1} \rightarrow \mathbb{k}^r, \quad p \mapsto (p(\xi_1), \dots, p(\xi_r)),$$

with respect to the monomial basis. As the rank of H_{d-1} is at most r , it follows that $\text{ev}_{\leq d-1}$ is surjective if and only if $\text{rk } H_{d-1} = r$.

As a computational reconstruction tool, Prony's method is summarized in Algorithm 1.1. If σ is a moment functional as before, then $m_\alpha = \sigma(x^\alpha)$, $\alpha \in \mathbb{N}^n$, denote the moments here.

Algorithm 1.1 Prony's method

Input: $r, d \in \mathbb{N}$ as well as $\{m_\alpha\}_{\alpha \in \mathbb{N}^n, |\alpha| \leq 2d}$.

Assumptions: There exist distinct points $\xi_1, \dots, \xi_r \in \mathbb{k}^n$ and parameters $\lambda_1, \dots, \lambda_r \in \mathbb{k} \setminus \{0\}$ from some field \mathbb{k} such that $m_\alpha = \sum_{j=1}^r \lambda_j \xi_j^\alpha$ for all $\alpha \in \mathbb{N}^n$, $|\alpha| \leq 2d$. Moreover, d is large enough such that $H_{d-1} := (m_{\alpha+\beta})_{|\alpha|, |\beta| \leq d-1}$ has rank r .

Output: λ_j, ξ_j for $1 \leq j \leq r$ (up to permutation).

- 1: Define $H_d := (m_{\alpha+\beta})_{|\alpha|, |\beta| \leq d}$.
 - 2: Compute $V(\ker H_d) = \{\xi_1, \dots, \xi_r\}$.
 - 3: Solve the linear system $m_\alpha = \sum_{j=1}^r \lambda_j \xi_j^\alpha$, $|\alpha| \leq 2d$, to obtain the coefficients $\lambda_1, \dots, \lambda_r$ corresponding to ξ_1, \dots, ξ_r .
-

Remark 1.3.2. Theorem 1.3.1 and Algorithm 1.1 are formulated in terms of the total degree of polynomials. Another common variant of Prony's method works with the max-degree of polynomials, which induces a filtration of the polynomial ring. More generally, one can work with other filtrations of the polynomial ring; see the statements in [vdOhe17, Chapter 2].

Another variation of Prony's method works with Toeplitz matrices of the form

$$\left(\sum_{j=1}^r \lambda_j \xi_j^{-\alpha+\beta} \right)_{|\alpha|_\infty, |\beta|_\infty \leq d}$$

instead of Hankel matrices, where the moments are usually bounded in max-degree. For this to be defined, the points ξ_1, \dots, ξ_r must have non-zero coordinates, so they are contained in the algebraic torus $(\mathbb{C}^*)^n$. This is especially common when working in a trigonometric setting, with points on the *complex torus*

$$\mathbb{T}^n := \{z \in \mathbb{C}^n \mid |z_1| = \dots = |z_n| = 1\}.$$

(Note that sometimes it is convenient to identify \mathbb{T}^n with the cubes $[-\pi, \pi]^n$ or $[0, 1]^n$ in \mathbb{R}^n , in which case we denote it by \mathbf{T}^n .) We will see more about this point of view in Chapters 3 and 4. \diamond

1.3.2 A multi-degree variant

In this section, we focus on a multi-degree version of Prony's method, which is useful, in particular, for the discussion of tensor decomposition in Section 1.5. We fix a field

\mathbb{k} and vector spaces $V_k := \mathbb{k}^{n_k}$ of dimension $n_k \in \mathbb{N}$, $1 \leq k \leq m$. Let $N := \mathbb{N}^m$ be an indexing set. Then, for $\alpha = (\alpha_1, \dots, \alpha_m) \in \mathbb{N}^{n_1} \times \dots \times \mathbb{N}^{n_m}$ and $\delta = (\delta_1, \dots, \delta_m) \in N$, we define $|\alpha| := (|\alpha_1|, \dots, |\alpha_m|) \in N$ and, using the partial order on N , we have $|\alpha| \leq \delta$ if $|\alpha_k| \leq \delta_k$ for $1 \leq k \leq m$.

Further, let $R_{\leq \delta} := \bigotimes_{k=1}^m \mathbb{k}[x_1, \dots, x_{n_k}]_{\leq \delta_k}$ and denote by $x^\alpha := x^{\alpha_1} \otimes \dots \otimes x^{\alpha_m}$, where $x^{\alpha_k} = x_1^{\alpha_{k1}} \dots x_{n_k}^{\alpha_{kn_k}}$ as usual. In other words, $R_{\leq \delta}$ is the polynomial ring in $\sum_{k=1}^m n_k$ variables truncated at multi-degree δ defined by a grading of the polynomial ring induced by N , so we can write $R_{\leq \delta} \cong \bigoplus_{|\alpha| \leq \delta} \mathbb{k}x^\alpha$.

Given points $\xi_{jk} \in V_k$, for $1 \leq j \leq r$, $1 \leq k \leq m$, we use the notation $\xi_j := (\xi_{j1}, \dots, \xi_{jm}) \in V_1 \times \dots \times V_m$ and $\xi_j^\alpha = \xi_{j1}^{\alpha_1} \dots \xi_{jm}^{\alpha_m}$. With this notation, the elements ξ_1, \dots, ξ_r are points in a space of dimension $\sum_{k=1}^m n_k$, the evaluation map corresponding to these points is given by

$$\text{ev}_{\leq \delta} := \text{ev}_{\leq \delta, \xi_1, \dots, \xi_r} : R_{\leq \delta} \longrightarrow \mathbb{k}^r, \quad x^\alpha \longmapsto (\xi_j^\alpha)_{1 \leq j \leq r}. \quad (1.2)$$

For ease of notation, we omit the points ξ_1, \dots, ξ_r from the index. The Vandermonde matrix corresponding to this map is

$$V_{\leq \delta} := \left(\xi_{j1}^{\alpha_1} \dots \xi_{jm}^{\alpha_m} \right)_{\substack{1 \leq j \leq r \\ |\alpha_k| \leq \delta_k, 1 \leq k \leq m}} = \left(\xi_j^\alpha \right)_{\substack{1 \leq j \leq r \\ |\alpha| \leq \delta}} \in \mathbb{k}^{r \times \binom{n_1 + \delta_1}{n_1} \dots \binom{n_m + \delta_m}{n_m}},$$

which we denote by $V_{\leq \delta}$ (not to be confused with the vector spaces V_k , $1 \leq k \leq m$).

Lemma 1.3.3 ([vdOhe17, Remark 2.8]). *Let $\sigma = \sum_{j=1}^r \lambda_j \text{ev}_{\xi_j}$ for $\lambda_j \in \mathbb{k}$ and $\xi_j \in V_1 \times \dots \times V_m$, $1 \leq j \leq r$. Let $\delta, \delta' \in N$. Then the following properties hold:*

- (1) $H_{\delta', \delta} := (\sigma(x^{\alpha+\beta}))_{|\alpha| \leq \delta', |\beta| \leq \delta} = V_{\leq \delta'}^\top \Lambda V_{\leq \delta}$, where $\Lambda := \text{diag}(\lambda_j)_{1 \leq j \leq r}$.
- (2) If $\lambda_1, \dots, \lambda_r \neq 0$ and $\text{ev}_{\leq \delta'} : R_{\leq \delta'} \rightarrow \mathbb{k}^r$, $x^\alpha \mapsto (\xi_j^\alpha)_{1 \leq j \leq r}$, is surjective, then $\ker V_{\leq \delta} = \ker H_{\delta', \delta}$.

Proof. The factorization $H_{\delta', \delta} = V_{\leq \delta'}^\top \Lambda V_{\leq \delta}$ follows by direct computation. Furthermore, if $\lambda_1, \dots, \lambda_r \neq 0$ and $\text{ev}_{\leq \delta'}$ is surjective, then $V_{\leq \delta'}^\top \Lambda$ represents an injective map, so the kernels of $V_{\leq \delta}$ and $H_{\delta', \delta}$ must agree. \square

Note that, if the points ξ_1, \dots, ξ_r are not distinct, then the map $\text{ev}_{\leq \delta'} : R_{\leq \delta'} \rightarrow \mathbb{k}^r$ can never be surjective, so the surjectivity assumption implies in particular that the points are distinct.

Lemma 1.3.4. *Let $\xi_j \in V_1 \times \dots \times V_m$, $1 \leq j \leq r$, let $\delta \in \mathbb{N}^m$ and $J \subseteq \{1, \dots, m\}$. Then the following are equivalent:*

- (1) $\text{im}(\text{ev}_{\leq \delta}) = \text{im}(\text{ev}_{\leq \delta'})$ for some $\delta' = \delta + \sum_{k \in J} a_k e_k$, $a_k \in \mathbb{Z}_{\geq 1}$,
- (2) $\text{im}(\text{ev}_{\leq \delta}) = \text{im}(\text{ev}_{\leq \delta'})$ for all $\delta' = \delta + \sum_{k \in J} a_k e_k$, $a_k \in \mathbb{Z}_{\geq 0}$,
- (3) $\text{im}(\text{ev}_{\leq \delta}) = \text{im}(\text{ev}_{\leq \delta + e_k})$ for all $k \in J$,

where $e_k \in \mathbb{N}^m$ denotes the k -th unit vector and $\text{ev}_{\leq \delta}$ denotes the truncated evaluation map (1.2) corresponding to ξ_1, \dots, ξ_r . Moreover, if $\text{im}(\text{ev}_{\leq \delta}) = \text{im}(\text{ev}_{\leq \delta - e_k})$ for all $k = 1, \dots, m$, then $\text{im}(\text{ev}_{\leq \delta}) = \text{im}(\text{ev}_{\xi_1, \dots, \xi_r})$.

Proof. First, observe that $\text{im}(\text{ev}_{\leq \delta}) \subseteq \text{im}(\text{ev}_{\leq \delta'})$ whenever $\delta \leq \delta'$. Therefore, if (1) holds, we have $\text{im}(\text{ev}_{\leq \delta}) = \text{im}(\text{ev}_{\leq \delta''}) = \text{im}(\text{ev}_{\leq \delta'})$ for all $\delta'' \in \mathbb{N}^m$ with $\delta \leq \delta'' \leq \delta'$. In particular, this holds for $\delta'' = \delta + e_k$, $k \in J$, so (3) is satisfied. The implication (2) \Rightarrow (1) is trivial. For the implication (3) \Rightarrow (2), we claim that $(\xi_j^\alpha)_{1 \leq j \leq r} \in \text{im}(\text{ev}_{\leq \delta})$ for all $|\alpha| \leq \delta'$. This follows by induction over $|\alpha|$. For this, we can assume that $|\alpha| \not\leq \delta$, so in particular α is non-zero. Choose $\tilde{\delta} \in \mathbb{N}^m$ such that $|\alpha| = \tilde{\delta} + e_k$ for some $k \in J$. This means, we can choose an element β such that $\beta < \alpha$ and $|\beta| = \tilde{\delta}$. Applying the inductive hypothesis, it follows that $(\xi_j^\beta)_j = \text{ev}_{\leq \delta}(p)$ for some $p \in R_{\leq \delta}$. Therefore, $(\xi_j^\alpha)_j = \text{ev}_{\leq \delta + e_k}(x_l p) \in \text{im}(\text{ev}_{\delta + e_k}) = \text{im}(\text{ev}_\delta)$, where l is the index in which α and β differ, which proves the claim.

The addendum follows by considering the case $J = \{k\}$ for $k = 1, \dots, m$ in order to conclude that $\text{im}(\text{ev}_{\leq \delta}) = \text{im}(\text{ev}_{\leq \delta + e_k})$ for all $k = 1, \dots, m$. \square

Lemma 1.3.5. *Let $\xi_j \in V_1 \times \dots \times V_m$, $1 \leq j \leq r$. If $\text{ev}_{\leq \delta - e_k}$ is surjective for all $k = 1, \dots, m$, then*

$$V(\ker \text{ev}_{\leq \delta}) = \{\xi_1, \dots, \xi_r\}.$$

Proof. This follows from [vdOhe17, Theorem 2.15]. \square

This yields the following multi-degree variant of Prony's method. It is a variant of [vdOhe17, Corollary 2.19] that we phrase in terms of a non-square Hankel matrix.

Theorem 1.3.6. *Let $\sigma = \sum_{j=1}^r \lambda_j \text{ev}_{\xi_j}$, where $\lambda_j \in \mathbb{k} \setminus \{0\}$, $\xi_j \in V_1 \times \dots \times V_m$ for $1 \leq j \leq r$. Let $\delta \in \mathbb{N}^m$ such that $\text{ev}_{\leq \delta - e_k}$ is surjective for all $k = 1, \dots, m$. Then*

$$V(\ker H_{\delta - e_k, \delta}) = \{\xi_1, \dots, \xi_r\}$$

for any choice of k , where $H_{\delta - e_k, \delta} := (\sigma(x^{\alpha + \beta}))_{|\alpha| \leq \delta - e_k, |\beta| \leq \delta}$.

Proof. First observe that the hypotheses imply that the points ξ_1, \dots, ξ_r are distinct. Then, by Lemma 1.3.3, we have $\ker V_{\leq \delta} = \ker H_{\delta - e_k, \delta}$ and it follows from Lemma 1.3.5 that $V(\ker H_{\delta - e_k, \delta}) = \{\xi_1, \dots, \xi_r\}$ for any $k \in \{1, \dots, m\}$. \square

Note that, by Lemma 1.3.4, the hypothesis that the evaluation maps $\text{ev}_{\leq \delta - e_k}$, $1 \leq k \leq m$, are surjective is satisfied if $\text{im}(\text{ev}_{\leq \delta}) = \text{im}(\text{ev}_{\leq \delta - e_k})$ for all $1 \leq k \leq m$ and the points ξ_1, \dots, ξ_r are distinct, as then $\text{im}(\text{ev}_{\leq \delta}) = \text{im}(\text{ev}_{\xi_1, \dots, \xi_r}) \cong \mathbb{k}^r$.

The case $m = 1$ of the theorem corresponds to the ordinary form of Prony's method with respect to the total degree of polynomials instead of multidegree – though, in contrast to Theorem 1.3.1, the Hankel matrix here is not square, which is a well-known optimization of Prony's method. Additionally, the following example shows that Theorem 1.3.6 gives a stronger statement than the variant of Prony's method that is formulated in Theorem 1.3.1 in terms of the total degree of polynomials.

Example 1.3.7. Let $m = 2$, $n_1 = n_2 = 1$, so that $V_1 \times V_2 \cong \mathbb{k}^2$. Moreover, let $\sigma = \lambda_1 \text{ev}_{\xi_1} + \lambda_2 \text{ev}_{\xi_2}$ be a map such that $\xi_{11} \neq \xi_{21}$ and $\xi_{12} \neq \xi_{22}$. Then

$$\text{im}(\text{ev}_{\leq(0,0)}) = \left\langle \left(\xi_1^{(0,0)}, \xi_2^{(0,0)} \right) \right\rangle = \langle (1, 1) \rangle$$

is a one-dimensional subspace of \mathbb{k}^2 , so $\text{ev}_{\leq(0,0)}$ is not surjective. However, both

$$\text{im}(\text{ev}_{\leq(1,0)}) = \left\langle (1, 1), \left(\xi_1^{(1,0)}, \xi_2^{(1,0)} \right) \right\rangle = \langle (1, 1), (\xi_{11}, \xi_{21}) \rangle$$

and

$$\text{im}(\text{ev}_{\leq(0,1)}) = \left\langle (1, 1), \left(\xi_1^{(0,1)}, \xi_2^{(0,1)} \right) \right\rangle = \langle (1, 1), (\xi_{12}, \xi_{22}) \rangle$$

are two-dimensional by choice of ξ_1, ξ_2 , so we can apply Theorem 1.3.6 with $\delta = (1, 1)$. Note that, in this example, both max-degree and total degree variants of Prony would require larger Vandermonde matrices, that is, those of max-degree 2 or total degree 2. Thus, Theorem 1.3.6 is a strictly stronger statement, in the sense that it requires fewer moments, i. e. evaluations of σ .

For the corresponding Hankel matrix, we either use moments of degree at most $(2, 1)$ or $(1, 2)$, instead of max-degree 4 or total degree 4. Thus, in total, this only requires 6 moments, whereas one would need 25 moments using max-degree or 15 moments using total degree. \diamond

Remark 1.3.8. A different kind of generalization is given by [Sau18]. For $r, n \in \mathbb{N}$, define the *hyperbolic cross*

$$\Upsilon_r := \left\{ \alpha \in \mathbb{N}^n \mid \prod_{i=1}^n (\alpha_i + 1) \leq r \right\},$$

whose cardinality is in $O(r \log^{n-1}(r))$, as follows from [PPST18, Example 8.21], for instance. Due to [Sau18, Corollary 11], it then holds that, for every set of r points $\xi_1, \dots, \xi_r \in \mathbb{C}^n$, the corresponding Vandermonde matrix $(\xi_j^\alpha)_{1 \leq j \leq r, \alpha \in \Upsilon_r}$ has rank r .

Define $\lceil \Upsilon_r \rceil$ as the Minkowski sum $\lceil \Upsilon_r \rceil := \Upsilon_r + \{0, e_1, \dots, e_n\}$. Then, for any $\sigma = \sum_{j=1}^r \lambda_j \text{ev}_{\xi_j}$, $\lambda_j \in \mathbb{C}^*$, $\xi_j \in \mathbb{C}^n$, $1 \leq j \leq r$, parameter recovery is possible by applying Prony's method to the Hankel matrix $(\sigma(x^{\alpha+\beta}))_{\alpha \in \Upsilon_r, \beta \in \lceil \Upsilon_r \rceil}$. The number of required moments is the cardinality of $\Upsilon_r + \Upsilon_r + \{0, e_1, \dots, e_n\}$, which is in $O(r^2 \log^{n-2}(r))$ for fixed $n \geq 2$, as follows from [Käm13, Corollary 4.9].

In Example 1.3.7 above, the hyperbolic cross $\Upsilon_2 = \{(0, 0), (0, 1), (1, 0)\}$ is of cardinality 3. The Hankel matrix involves the evaluations at exponents from the set $\Upsilon_2 + \lceil \Upsilon_2 \rceil = \{\alpha \in \mathbb{N}^2 \mid |\alpha| \leq 3\}$, which has cardinality 10. Note that this is larger than the number of moments needed in Example 1.3.7 because the statement here holds universally for any set of points, whereas, in the example, the choice of δ depends on the concrete instance of a point set. \diamond

1.4 Special case of equal weights

An interesting special case of the reconstruction problem described in Section 1.3.1 arises when the given moments correspond to evaluations of a univariate exponential sum for which all the weights are equal. For simplicity, we can assume that all the weights are 1. In this case, we have $n = 1$ and the moments are of the form

$$m_k = \xi_1^k + \cdots + \xi_r^k, \quad k \in \mathbb{N},$$

for (not necessarily distinct) points $\xi_j \in \mathbb{k}$, $1 \leq j \leq r$, where we assume that \mathbb{k} is a field of characteristic 0. This means that the moments are symmetric polynomials in the a priori unknown variables ξ_j , as they are invariant under permutations of the points ξ_1, \dots, ξ_r . More specifically, these are exactly the *power sums* ([Mac95, Chapter 1.2]), which form a basis of the ring of symmetric functions in r variables, so that every symmetric polynomial in the unknowns ξ_j can be expressed formally in terms of the moments m_k , $k \in \mathbb{N}$. Since the basis is compatible with the degree, any symmetric polynomial of degree $d \in \mathbb{N}$ can in fact be expressed in terms of m_0, \dots, m_d .

A different basis of the ring of symmetric polynomials consists of the *elementary symmetric polynomials* ([Mac95, Chapter 1.2])

$$e_k := \sum_{1 \leq j_1 < \cdots < j_k \leq r} \xi_{j_1} \cdots \xi_{j_k}, \quad k \in \mathbb{N}.$$

These have the property

$$\prod_{j=1}^r (X - \xi_j) = e_0 X^r - e_1 X^{r-1} + \cdots + (-1)^r e_r.$$

In other words, the coefficients of the Prony polynomial $\prod_{j=1}^r (X - \xi_j)$ are given, up to sign, by the elementary symmetric polynomials in terms of the points ξ_j .

As both the power sums and the elementary symmetric polynomials generate the ring of symmetric polynomials, the elementary symmetric polynomials can be computed from the power sums by a non-linear change of basis; see [Mac95, Chapter 1.2, Equation (2.14')]. For the parameter recovery problem, this means that we can compute the $r + 1$ coefficients of the Prony polynomial from the moments m_0, \dots, m_r , by this change of basis.

Subsequently computing the roots of the Prony polynomial allows to recover the points ξ_1, \dots, ξ_r . If the points ξ_1, \dots, ξ_r are not known to be distinct, the roots need to be computed with multiplicity. In summary, we obtain the following statement.

Proposition 1.4.1. *Let m_k , $k \in \mathbb{N}$, be moments of the form $m_k = \xi_1^k + \cdots + \xi_r^k$, where $\xi_1, \dots, \xi_r \in \mathbb{k}$. Then the parameters ξ_1, \dots, ξ_r can be recovered uniquely from the moments m_0, \dots, m_r .*

Proof. This follows from the above discussion. Note in particular that $m_0 = r$ yields the number of summands. \square

This approach requires fewer moments than Prony's method in Theorem 1.3.1 and Algorithm 1.1 would use; namely, $r + 1$ instead of $2r + 1$. However, this involves a non-linear transformation for the change of basis, whereas Prony's method would require solving a linear problem for the computation of the kernel of a moment matrix.

Example 1.4.2. Let $r = 4$ and consider the moment vector given by $(m_0, \dots, m_4) = (4, 14, 50, 182, 674)$. We want to construct the univariate Prony polynomial

$$f := e_0 X^4 - e_1 X^3 + e_2 X^2 - e_3 X + e_4$$

whose coefficients are elementary symmetric polynomials. In terms of the power sums, we have

$$f = p_0 X^4 - p_1 X^3 + \left(\frac{1}{2}p_{1,1} - \frac{1}{2}p_2\right) X^2 + \left(-\frac{1}{6}p_{1,1,1} + \frac{1}{2}p_{2,1} - \frac{1}{3}p_3\right) X + \frac{1}{24}p_{1,1,1,1} - \frac{1}{4}p_{2,1,1} + \frac{1}{8}p_{2,2} + \frac{1}{3}p_{3,1} - \frac{1}{4}p_4,$$

where the coefficients p_{k_1, \dots, k_l} denote the power sum basis elements, which are indexed by partitions. Since $p_{k_1, \dots, k_l} = m_{k_1} \cdots m_{k_l}$, we obtain

$$f = X^4 - 14X^3 + 73X^2 - 168X + 144 = (X - 3)^2(X - 4)^2.$$

Thus, the factorization of f shows that the moments are of the form $m_k = 2 \cdot 3^k + 2 \cdot 4^k$ for $0 \leq k \leq 4$. In particular this example shows that the points ξ_1, \dots, ξ_r do not need to be distinct. We can count the roots with multiplicities which allows for weights that are different from 1, but are (small) positive natural numbers instead. \diamond

For related works and generalizations, see for example [Tsa+20] and [MSW21].

1.5 Connection to symmetric tensor decomposition

Here, we briefly explore the relationship between Prony's method and the decomposition of tensors. For extensive references on this topic, we refer to [Lan12] as well as [IK99], but also [BCMT10; BBCM13; Mou18] are relevant. We start with a short introduction to tensors.

Let \mathbb{k} be a field of characteristic 0. Let $V = \mathbb{k}^{n+1}$, for some $n \in \mathbb{N}$, be a \mathbb{k} -vector space and let $d \in \mathbb{N}$. The d -fold tensor product $V^{\otimes d}$ is again a vector space whose elements are called *tensors of order d* . Tensors of the form $v_1 \otimes \cdots \otimes v_d \in V^{\otimes d}$ are called *elementary tensors* or tensors of rank 1, if $v_1, \dots, v_d \in V \setminus \{0\}$. As $V^{\otimes d}$ is finite-dimensional, namely of dimension $(n + 1)^d$, each tensor $t \in V^{\otimes d}$ can be written as a finite sum of elementary tensors

$$t = \sum_{j=1}^r v_{j1} \otimes \cdots \otimes v_{jd}, \tag{1.3}$$

where $v_{jk} \in V$ for $1 \leq j \leq r$, $1 \leq k \leq d$. If r is the smallest integer for which this is possible, then r is called *rank* of t and (1.3) is a (rank-) *decomposition* of the tensor t (also called *canonical-polyadic decomposition*). In general, it is difficult to compute the rank of a tensor or a decomposition of a tensor if the order d of the tensor is larger than 2. Tensors of order 2 correspond to matrices, for which efficient rank computations are available.

The space of *symmetric tensors* $S^d(V)$ is defined as the $\binom{n+d}{d}$ -dimensional subspace of $V^{\otimes d}$ of tensors that are invariant under the permutation action of the symmetric group S_d defined by $v_1 \otimes \cdots \otimes v_d \mapsto v_{\pi(1)} \otimes \cdots \otimes v_{\pi(d)}$ for $\pi \in S_d$. Every symmetric tensor $t \in S^d(V)$ can be written as

$$t = \sum_{j=1}^r \lambda_j v_j^{\otimes d} \quad (1.4)$$

with suitable vectors $v_j \in V$ and coefficients $\lambda_j \in \mathbb{k}$, $1 \leq j \leq r$. The minimal possible r is called *symmetric rank* of t and (1.4) is a symmetric(-rank) decomposition. Note that the symmetric rank of a symmetric tensor can be strictly larger than its rank when viewed as a tensor in $V^{\otimes d}$; see [Shi18] for an example. If the field \mathbb{k} is algebraically closed, the coefficients λ_j in the decomposition (1.4) can be omitted, as then $t = \sum_{j=1}^r w_j^{\otimes d}$, where $w_j := \sqrt[d]{\lambda_j} v_j \in V$ for $1 \leq j \leq r$.

It is possible to identify the set of symmetric tensors $S^d(V)$ with $\mathbb{k}[X_0, \dots, X_n]_d$, the set of homogeneous polynomials of degree d in $n+1$ variables. If $x_0, \dots, x_n \in V$ is a basis of V , then, in terms of the basis of $V^{\otimes d}$, the identification is defined by the map

$$S^d(V) \longrightarrow \mathbb{k}[X_0, \dots, X_n]_d, \\ \sum_{\beta \in \mathbb{N}^d, |\beta|_\infty \leq n} t_\beta x_{\beta_1} \otimes \cdots \otimes x_{\beta_d} \longmapsto \sum_{\beta} t_\beta X_{\beta_1} \cdots X_{\beta_d} = \sum_{\alpha \in \mathbb{N}^{n+1}, |\alpha|=d} \binom{d}{\alpha} \hat{t}_\alpha X^\alpha, \quad (1.5)$$

where $\hat{t}_\alpha = t_\beta$ if $X_{\beta_1} \cdots X_{\beta_d} = X^\alpha$, which is well-defined for symmetric tensors. Equivalently, the correspondence is given by

$$\sum_{j=1}^r \lambda_j v_j^{\otimes d} \longmapsto \sum_{j=1}^r \lambda_j L_j^d,$$

where $v_j = \sum_{i=0}^n v_{ji} x_i$ and $L_j = \sum_{i=0}^n v_{ji} X_i$ with coefficients $v_{ji} \in \mathbb{k}$ for $1 \leq j \leq r$, $0 \leq i \leq n$. Thus, finding a symmetric decomposition of a tensor is equivalent to decomposing the corresponding homogeneous degree- d polynomial as a linear combination of powers of linear forms with the smallest possible number of summands. This is related to a classical problem, known as *Waring problem* (see for instance [Lan12, Chapter 5.4]).

Remark 1.5.1. The occurrence of multinomial coefficients $\binom{d}{\alpha}$ in the map (1.5) illustrates why, for simplicity of exposition, we work over a field of characteristic 0. The identification is also possible whenever the characteristic of the field is larger than d . More generally, one can work with divided powers instead, which is a concept we do not discuss here. See [IK99] for an extensive treatment of this view point. \diamond

As the rank of a tensor does not change under scalar multiplication of a tensor, it is often convenient to work with the projective spaces $\mathbb{P}S^d(V) \cong \mathbb{P}_{\mathbb{k}}^{\binom{n+d}{d}-1}$ and $\mathbb{P}(V) \cong \mathbb{P}_{\mathbb{k}}^n =: \mathbb{P}^n$. In light of this, the following definition is illustrative.

Definition 1.5.2. The d -uple Veronese embedding of \mathbb{P}^n is given by

$$\nu_d: \mathbb{P}^n \longrightarrow \mathbb{P}^{\binom{n+d}{d}-1}, \quad [x_0 : \cdots : x_n] \longmapsto [x^\alpha]_{|\alpha|=d}.$$

The image $\nu_d(\mathbb{P}^n)$ is called *Veronese variety*.

Thus, the Veronese embedding projectively embeds vectors from V into the set of symmetric d -tensors (up to a projective isomorphism that scales the coordinates by $\binom{d}{\alpha}$). The Veronese variety is a projective variety in $\mathbb{P}^{\binom{n+d}{d}-1}$ that parametrizes exactly the projective classes of symmetric rank-1 tensors (again up to projective isomorphism).

On the algebraic side, the Veronese embedding ν_d corresponds to the ring homomorphism

$$\begin{aligned} \mathbb{k}[m_\alpha \mid \alpha \in \mathbb{N}^n, |\alpha| \leq d] &\longrightarrow \mathbb{k}[X_0, \dots, X_n], \\ m_\alpha &\longmapsto X_0^{d-|\alpha|} X_1^{\alpha_1} \cdots X_n^{\alpha_n}, \end{aligned} \tag{1.6}$$

whose image is the subalgebra of the polynomial ring generated by all monomials of degree d . The kernel of this homomorphism is the defining ideal of the Veronese variety in $\mathbb{P}^{\binom{n+d}{d}-1}$. The reason for our peculiar asymmetric choice of $\{\alpha \in \mathbb{N}^n, |\alpha| \leq d\}$ as indexing set for the variables m_α is that, with this notation, generators for the ideal defining the Veronese variety are given by the 2×2 -minors of the (generalized) Hankel matrix

$$(m_{\alpha+\beta})_{|\alpha| \leq 1, |\beta| \leq d-1}, \tag{1.7}$$

so the ideal is generated in degree 2 and has a determinantal representation (cf. [Har87, Exercise 1.2.12], [Lan12, Section 6.10.4]). This matrix is of size $\dim V \times \dim S^{d-1}(V) = (n+1) \times \binom{n+d-1}{d-1}$. In fact, it is induced by the polarization map

$$S^d(V) \longrightarrow S^1(V) \otimes S^{d-1}(V), \quad v^{\otimes d} \longmapsto v \otimes v^{\otimes(d-1)},$$

which views a symmetric tensor as a *partially-symmetric* tensor; in this case, a tensor that is invariant under permutations of the last $d-1$ factors in $V^{\otimes d}$. This is called a *flattening* of the symmetric tensor. More generally, one can consider flattenings of the form $S^d(V) \rightarrow S^k(V) \otimes S^{d-k}(V)$ for $0 \leq k \leq d$ (also called $(k, d-k)$ -flattenings), which also give rise to some vanishing equations (cf. [Lan12, Section 3.5]).

Example 1.5.3. Let $x, y \in \mathbb{k}^2$ be a basis and consider the symmetric 3-tensor

$$\lambda \cdot x \otimes x \otimes x + \mu \cdot (x \otimes x \otimes y + x \otimes y \otimes x + y \otimes x \otimes x),$$

where $\lambda, \mu \in \mathbb{k}$. In terms of its slices, the tensor could also be written as

$$\left(\begin{array}{cc|cc} \lambda & \mu & \mu & 0 \\ \mu & 0 & 0 & 0 \end{array} \right).$$

Under the mapping (1.5), this tensor corresponds to the binary 3-form

$$\lambda \cdot X^3 + \mu \cdot \binom{3}{2,1} X^2 Y = \lambda \cdot X^3 + \mu \cdot 3X^2 Y.$$

The (1, 2)-flattening is

$$\lambda \cdot x \otimes x^2 + \mu \cdot x \otimes 2xy + \mu \cdot y \otimes x^2,$$

which is the same as the Hankel matrix

$$\begin{pmatrix} \lambda & \mu & 0 \\ \mu & 0 & 0 \end{pmatrix}$$

in terms of the bases x, y and $x^2, 2xy, y^2$ of \mathbb{k}^2 and $S^2(\mathbb{k}^2)$, respectively. If $\mu \neq 0$, then the 2×2 -minors of this matrix do not all vanish, which shows that the tensor is not of symmetric rank 1, but the symmetric rank must be larger. On the other hand, if $\mu = 0$, then $\lambda x \otimes x \otimes x$ is clearly of symmetric rank ≤ 1 . \diamond

As mentioned above, the Veronese variety parametrizes projective classes of symmetric rank-1 tensors. A tensor of symmetric rank 2 is of the form $t = \lambda_1 \xi_1^{\otimes d} + \lambda_2 \xi_2^{\otimes d}$ for $\xi_1, \xi_2 \in V$ with $\xi_1 \neq \xi_2$ and $\lambda_1, \lambda_2 \in \mathbb{k}^*$. Then $[t] \in \mathbb{P}(V)$ is contained in the projective line spanned by the points $[\xi_1^{\otimes d}], [\xi_2^{\otimes d}] \in \mathbb{P}(V)$, a secant line to the Veronese variety. The Zariski closure of the union of all secant lines forms the (*second*) *secant variety* of the Veronese variety, denoted by $\sigma_2(\nu_d(\mathbb{P}^n))$. It is not parametrized by the set of projective classes of symmetric rank-2 tensors, but these tensors form a Zariski-dense subset of this secant variety. This means that not every point on the secant variety corresponds to a tensor of symmetric rank 2, but some points correspond to tensors of possibly larger symmetric rank – such tensors are said to have symmetric *border rank* 2.

More generally, one can define the r -th secant variety $\sigma_r(\nu_d(\mathbb{P}^n))$, which is the Zariski closure of the classes of symmetric rank- r tensors. The projective point $[\sum_{j=1}^r \lambda_j \xi_j^{\otimes d}]$ for $\lambda_j \in \mathbb{k}$, $\xi_j \in V$, $1 \leq j \leq r$, lies in the $(r-1)$ -dimensional projective subspace spanned by $\nu_d(\xi_1), \dots, \nu_d(\xi_r)$. The $(r+1) \times (r+1)$ -minors of flattenings of symmetric tensors of order d give rise to some equations of the r -th secant variety of the Veronese variety, but often these are not enough to generate the defining ideal and in general it is hard to find generators of this ideal. We refer to [Lan12, Chapter 7] for an introduction to this topic. This is also studied in detail in [IK99] in terms of the catalecticant matrix.

Similarly, we may consider general (non-symmetric) order- d tensors in $V_1 \otimes \dots \otimes V_d$, where V_1, \dots, V_d are \mathbb{k} -vector spaces of dimensions $n_k + 1$ for some $n_k \in \mathbb{N}$, $1 \leq k \leq d$. These tensors can be decomposed as

$$\sum_{j=1}^r v_{j1} \otimes \dots \otimes v_{jd},$$

with $v_{jk} \in V_k \setminus \{0\}$ for $1 \leq j \leq r$, $0 \leq k \leq d$ and some $r \in \mathbb{N}$. Moreover, they can be associated with polynomials in N variables that are homogeneous with respect to a suitable multidegree. With the identification $\mathbb{P}(V_k) \cong \mathbb{P}^{n_k}$, the following definition is relevant.

Definition 1.5.4. The *Segre embedding* of $\mathbb{P}^{n_1} \times \cdots \times \mathbb{P}^{n_d}$ is defined as

$$\begin{aligned} \mathbb{P}^{n_1} \times \cdots \times \mathbb{P}^{n_d} &\longrightarrow \mathbb{P}^{N-1}, \\ ([x_{1,i_1}]_{i_1}, \dots, [x_{d,i_d}]_{i_d}) &\longmapsto [x_{1,i_1} \cdots x_{d,i_d}]_{i_1, \dots, i_d}, \end{aligned}$$

where $N = \prod_{k=1}^d (n_k + 1)$ and $i_k \in \{0, \dots, n_k\}$ for $1 \leq k \leq d$. The image of this map is called *Segre variety*.

If $v_k \in V_k$, $1 \leq k \leq d$, are vectors with coordinates $v_k = (x_{k,0}, \dots, x_{k,n_k})$, then the rank-1 tensor $v_1 \otimes \cdots \otimes v_d$ has coordinates $(x_{1,i_1} \cdots x_{d,i_d})_{i_1, \dots, i_d}$ whose projective class is the image of $([v_1], \dots, [v_d])$ under the Segre embedding. As such, the Segre variety is parametrized by projective classes of rank-1 tensors in $\mathbb{P}(V_1 \otimes \cdots \otimes V_d) \cong \mathbb{P}^{N-1}$. Similarly, for $r \geq 2$, classes of rank- r tensors are contained in the r -th secant variety of the Segre variety, but do not in general fully parametrize it.

More generally, one can consider partially-symmetric tensors in $S^{\delta_1}(V_1) \otimes \cdots \otimes S^{\delta_d}(V_d)$, which are tensors of order $\delta_1 + \cdots + \delta_d$, where $\delta_1, \dots, \delta_d \in \mathbb{N}$. These correspond to polynomials in N variables that are homogeneous of multidegree $(\delta_1, \dots, \delta_d)$, for a suitable choice of multidegree. Analogously, one defines decompositions of partially-symmetric tensors and the corresponding *Segre-Veronese* embeddings and varieties.

Remark 1.5.5. After this brief introduction to tensors, let us more closely investigate the relationship to Prony's method. For this, we consider the ring homomorphism (1.6) again. In contrast to the tensor setting before, we now work with inhomogeneous coordinates. Let $S := \mathbb{k}[m_\alpha \mid \alpha \in \mathbb{N}^n, |\alpha| \leq d]$ be the domain of the map (1.6) and let $R := \mathbb{k}[x_1, \dots, x_n]$.

Given any \mathbb{k} -linear map $\sigma: R \rightarrow \mathbb{k}$, we may interpret the coordinates of the space $\mathbb{P}^{\binom{n+d}{d}-1}$, the variables m_α , as representations of the moments $\sigma(x^\alpha)$ for $|\alpha| \leq d$. In other words, the restricted map $\sigma|_{R_{\leq d}}$ factors through the evaluation homomorphism $S \rightarrow \mathbb{k}$, defined by $m_\alpha \mapsto \sigma(x^\alpha)$, and the \mathbb{k} -linear map $R_{\leq d} \rightarrow S$, $x^\alpha \mapsto m_\alpha$.

In particular, if σ is of the form $\sigma(x^\alpha) = \xi^\alpha$, $|\alpha| \leq d$, for some $\xi \in \mathbb{k}^n$ – the moment functional of the Dirac measure δ_ξ supported at the point ξ – then all the 2×2 -minors of the Hankel moment matrix

$$(\sigma(x^{\alpha+\beta}))_{|\alpha| \leq 1, |\beta| \leq d-1}$$

vanish. As this is the image of the matrix (1.7) under the evaluation homomorphism, the moment vector $[\sigma(x^\alpha)]_{|\alpha| \leq d} \in \mathbb{P}^{\binom{n+d}{d}-1}$ corresponds to a point on the Veronese variety, namely the point $\nu_d([1 : \xi_1 : \cdots : \xi_n])$.

More generally, if σ is of the form $\sigma(x^\alpha) = \sum_{j=1}^r \lambda_j \xi_j^\alpha$ for $\lambda_j \in \mathbb{k}$, $\xi_j \in \mathbb{k}^n$ – the moment functional of a signed measure supported at r points – then the moments correspond to points on the r -th secant variety. Put differently, the moment vector $(\sigma(x^\alpha))_{|\alpha| \leq d}$ corresponds to a symmetric tensor of symmetric rank at most r , as we may associate it

with the homogeneous degree- d polynomial

$$\sum_{j=1}^r \lambda_j (X_0 + \xi_{j1}X_1 + \cdots + \xi_{jn}X_n)^d = \sum_{|\alpha| \leq d} \binom{d}{|\alpha|, \alpha} \sigma(x^\alpha) X_0^{d-|\alpha|} X_1^{\alpha_1} \cdots X_n^{\alpha_n}. \quad (1.8)$$

Now it is apparent that recovery of the parameters λ_j, ξ_j of σ from its moments up to degree d also produces a decomposition of the corresponding symmetric tensor into a sum of at most r symmetric rank-1 tensors. Vice versa, if the symmetric tensor is of symmetric rank r and has a unique decomposition, then this decomposition also allows us to retrieve the parameters λ_j, ξ_j for $1 \leq j \leq r$ that define σ . The property of having a unique decomposition is also referred to as *identifiability*. \diamond

Example 1.5.6. Let $d, r \in \mathbb{N}$ with $d \geq 1$ and $r \leq \binom{n+d-1}{d-1}$ and let $V = \mathbb{C}^{n+1}$ for some $n \in \mathbb{N}$, $n \geq 1$. Let $t \in S^{2d-1}(V)$ be a generic symmetric tensor of symmetric rank r . Then the decomposition can be computed via Prony's method by Theorem 1.3.6.

For this, first note that the case $d = 1$ is trivial as then $t \in V$. Thus, we can assume that $d \geq 2$, so t is a tensor of order at least 3. The case $n = 1$ corresponds to binary forms which can be handled separately; classically this is dealt with by Sylvester's algorithm, which is a homogeneous formulation of the classic univariate Prony method and goes back to [Sy186]. In this case, we have $r \leq d$ and t has a unique decomposition by [IK99, Theorem 1.40]. When $n \geq 2$, one checks explicitly that the tensor is of subgeneric symmetric rank, i. e. it holds that $r < \frac{1}{n+1} \binom{n+2d-1}{2d-1}$, which follows from our choice of r . By [COV17], a generic symmetric tensor of subgeneric symmetric rank has a unique decomposition, unless it belongs to a small number of exceptional cases. The exceptional cases do not apply in this example, so we conclude that the tensor t can be decomposed uniquely, as it is generic among the symmetric tensors of symmetric rank r .

Now that we have established the uniqueness of the decomposition, let us focus on the recovery of parameters. Without loss of generality, we can assume that the tensor is of the form $t = \sum_{j=1}^r \lambda_j (1, \xi_{j1}, \dots, \xi_{jn})^{\otimes (2d-1)}$ for $\lambda_j \in \mathbb{C}^*$ and distinct elements $\xi_j \in \mathbb{C}^n$, $1 \leq j \leq r$, as t was assumed to be generic. By the correspondence (1.8), we can view the entries of t as evaluations $\sigma(x^\alpha)$, $|\alpha| \leq 2d-1$, of the functional

$$\sigma: \mathbb{C}[x_1, \dots, x_n] \longrightarrow \mathbb{C}, \quad x^\alpha \longmapsto \sum_{j=1}^r \lambda_j \xi_j^\alpha.$$

From this, we construct the Hankel moment matrix $\sigma(x^{\alpha+\beta})_{|\alpha| \leq d-1, |\beta| \leq d}$. It can also be obtained as the $(d-1, d)$ -flattening of t .

In order to apply Theorem 1.3.6, we need to assert that the Vandermonde matrix $(\xi_j^\alpha)_{|\alpha| \leq d-1, 1 \leq j \leq r}$ has rank r . As the points ξ_1, \dots, ξ_r are generic and since we have assumed that $r \leq \binom{n+d-1}{d-1}$, this requirement is satisfied, as follows from apolarity theory by [IK99, Lemma 1.15]. Thus, finally, we are able to apply Theorem 1.3.6 with $m = 1$ to retrieve the points ξ_1, \dots, ξ_r and subsequently the weights $\lambda_1, \dots, \lambda_r$, by solving a linear system. \diamond

Remark 1.5.7. Note that the cases in which Prony’s method can successfully decompose a symmetric tensor, as in Example 1.5.6, are, in a sense, the less interesting cases. These correspond to cases in which flattenings give rise to equations for the r -th secant variety of the Veronese variety. The requirement $r \leq \binom{n+d-1}{d-1}$ means that the symmetric rank r must be quite small compared to the order of the tensor, $2d - 1$, which is not a typical scenario in many applications of tensor decompositions. For larger ranks, it is much more challenging to find decompositions and defining equations – in particular for tensors of *generic* rank – and the search for better tools for these is an active area of research, with a large variety of available results for specific small values of n , d and r . A direct adaption of Prony’s method or Sylvester’s algorithm to the decomposition of symmetric tensors of higher rank is possible using a symbolic approach by extending the Hankel matrix in a rank-preserving way, as proposed in [BCMT10].

More generally, one can apply Theorem 1.3.6 to certain partially-symmetric tensors by considering flattenings of the form

$$S^{2\delta_1-1}(V_1) \otimes \bigotimes_{k=2}^m S^{2\delta_k}(V_k) \longrightarrow \left(S^{\delta_1-1}(V_1) \otimes \bigotimes_{k=2}^m S^{\delta_k}(V_k) \right) \otimes \left(\bigotimes_{k=1}^m S^{\delta_k}(V_k) \right),$$

for some $\delta := (\delta_1, \dots, \delta_m) \in \mathbb{N}^m$ and vector spaces V_1, \dots, V_m , as elements in the space on the right can be viewed as Hankel matrices whose entries have the multi-indices $\alpha + \beta$, $|\alpha| \leq \delta - e_1$, $|\beta| \leq \delta$. For the assumptions of Theorem 1.3.6 to be satisfied, the partially-symmetric rank of the tensor must be sufficiently small, but we omit a more detailed analysis of the requirements here.

Although the above does not cover general tensors without symmetries, there are generalizations for the decomposition of general tensors such as [BBCM13], as well. Later, in Example 4.3.16, we will also see more explicitly how to find a decomposition of a general third-order tensor of small rank through the computation of generalized eigenvalues of the slices of the tensor. \diamond

Example 1.5.8. A frequent scenario consists of exponential sums or discrete signed measures on the torus. For this, assume that the points ξ_1, \dots, ξ_r lie in $(\mathbb{C}^*)^m$, the m -dimensional algebraic torus over \mathbb{C} . Let $\sigma: \mathbb{C}[x_1^{\pm 1}, \dots, x_m^{\pm 1}] \rightarrow \mathbb{C}$ be a \mathbb{C} -linear map from the Laurent polynomial ring of the form $\sigma(x^\alpha) = \sum_{j=1}^r \lambda_j \xi_j^\alpha$ for some $\lambda_j \in \mathbb{C}^*$ and consider its moments $\sigma(x^\alpha)$ for $-\delta_i \leq \alpha_i \leq \delta_i$, $1 \leq i \leq m$, and some $\delta_1, \dots, \delta_m \in \mathbb{N}$, i. e. the moments are chosen symmetrically around zero in each direction. By a translation on the torus, we have

$$\sigma(x^\alpha) = \sum_{j=1}^r \lambda_j \xi_j^\alpha = \sum_{j=1}^r \lambda_j \xi_j^{-\delta} \xi_j^{\alpha+\delta},$$

where $\delta := (\delta_1, \dots, \delta_m)$. As the exponents involving α on the right are non-negative, the problem can be associated to the partially-symmetric tensor decomposition problem in $S^{2\delta_1}(\mathbb{C}^2) \otimes \dots \otimes S^{2\delta_m}(\mathbb{C}^2)$ corresponding to the tensor with weights $\lambda_j \xi_j^{-\delta}$ and rank-1 tensors

$$(1, \xi_{j1})^{\otimes 2\delta_1} \otimes \dots \otimes (1, \xi_{jm})^{\otimes 2\delta_m},$$

for $1 \leq j \leq r$. ◇

In the context of tensor decomposition, one is often interested in finding a low-rank approximation of a tensor. In order to quantify how well a low-rank tensor approximates a given tensor, one needs to choose a norm on the space of tensors. The following example explores what a reasonable choice of norm could be for symmetric tensors – a norm that does not depend on the choice of coordinates of the underlying space.

Example 1.5.9. We consider a symmetric rank-1 tensor $v^{\otimes d''}$ for some vector $v \in V := \mathbb{C}^{n+1}$, $n \in \mathbb{N}$, as well as an associated flattening of the form

$$v^{\otimes d''} = v^{\otimes d} \otimes v^{\otimes d'},$$

where $d'' = d + d'$ and $d, d' \in \mathbb{N}$. The flattening of the tensor can be viewed as a partially symmetric tensor in $S^d(V) \otimes S^{d'}(V) = U \otimes W$.

We define a norm on the space $U = S^d(V)$ in terms of the *Bombieri-2-norm* $[-]_2$ of homogeneous polynomials (cf. [BBEM90]). For a homogeneous polynomial $f = \sum_{|\alpha|=d} f_\alpha X^\alpha \in \mathbb{C}[X_0, \dots, X_n]$ of degree d , it is defined by

$$[f]_2 := \left(\sum_{|\alpha|=d} \frac{|f_\alpha|^2}{\binom{d}{\alpha}} \right)^{\frac{1}{2}}.$$

By the correspondence (1.5), we obtain for a symmetric tensor $t \in S^d(V)$ that

$$[t]_2 = \left(\sum_{|\alpha|=d} \binom{d}{\alpha} |\hat{t}_\alpha|^2 \right)^{\frac{1}{2}},$$

where $\hat{t}_\alpha = \binom{d}{\alpha}^{-1} f_\alpha$ are the coefficients of the symmetric tensor t as defined in (1.5). Thus, for the elementary tensor $v^{\otimes d} \in S^d(V)$, we have

$$[v^{\otimes d}]_2 = \left(\sum_{|\alpha|=d} \binom{d}{\alpha} |v^\alpha|^2 \right)^{\frac{1}{2}} = \left(\sum_{i=0}^n |v_i|^2 \right)^{\frac{d}{2}} = \|v\|_2^d.$$

More generally, for arbitrary symmetric tensors, this norm is invariant under isometries of the underlying vector space V , which makes it a natural choice. A proof of this property is given in [BC13, Theorem 16.3].

On the component $W = S^{d'}(V)$, we define the norm in an analogous way. Finally, we choose a norm $\|-\|$ on the space $U \otimes W$ such that rank-1 elements satisfy

$$\|u \otimes w\| = [u]_2 \cdot [w]_2 \tag{1.9}$$

and therefore $\|v^{\otimes d} \otimes v^{\otimes d'}\| = \|v\|_2^{d+d'}$. A norm that satisfies (1.9) is a *cross-norm*. For instance, this property holds for the Frobenius norm and the spectral norm, but also for the induced projective norm (cf. [Hac19, Lemma 4.51]). ◇

For further properties of norms on tensor spaces, we refer to [Hac19, Chapter 4.2].

2 Local Dirac mixtures

The contents of this chapter are adapted from our article [GW20]. Unless stated otherwise, we assume that \mathbb{k} denotes an algebraically closed field of characteristic 0.

Local mixture distributions are distributions that play a role in some statistical settings; see [Mar02; AM07]. They involve a distribution ϕ_ξ and its derivatives $\phi_\xi^{(i)}$ and are of the form

$$\phi_\xi(x) + \sum_{i=1}^l \lambda_i \phi_\xi^{(i)}(x),$$

where $\xi, \lambda_i, 1 \leq i \leq l$, are some parameters. We call this a *local mixture distribution* of order l .

In this chapter, we primarily focus on the degenerate case of univariate Dirac distributions, so we consider distributions of the form

$$\delta_\xi + \lambda_1 \delta'_\xi + \cdots + \lambda_l \delta_\xi^{(l)},$$

for $\xi, \lambda_i \in \mathbb{k}, 1 \leq i \leq l$, where $\delta'_\xi, \dots, \delta_\xi^{(l)}$ denote the distributional derivatives of the Dirac distribution δ_ξ ; see [Sch73, Chapter 2]. Note that, while the Dirac distribution δ_ξ is also a (non-negative) measure, its derivatives are not non-negative or signed measures. As the moments of the Dirac distribution are of a particularly simple form, the moments of a local mixture of a Dirac distribution can be computed easily, as well. More precisely, the k -th moment of a local mixture of a Dirac distribution δ_ξ of order l is of the form

$$m_k = \xi^k + \sum_{i=1}^{\min\{l, k\}} \lambda_i \frac{k!}{(k-i)!} \xi^{k-i},$$

for $k \in \mathbb{N}$ and certain parameters $\xi, \lambda_1, \dots, \lambda_l \in \mathbb{k}$ (cf. [Sch73, Chapter 2]). These are algebraic expressions in the parameters $\xi, \lambda_i, 1 \leq i \leq l$, which allows us to study these distributions from an algebraic point of view.

More generally, we consider mixtures of such distributions centered at different points, that is, distributions of the form

$$\sum_{j=1}^r \lambda_{j0} \delta_{\xi_j} + \lambda_{j1} \delta'_{\xi_j} + \cdots + \lambda_{j, l_j} \delta_{\xi_j}^{(l_j)},$$

for r points $\xi_1, \dots, \xi_r \in \mathbb{k}$ and weights $\lambda_{j,k_j} \in \mathbb{k}$, for $0 \leq k_j \leq l_j$, $1 \leq j \leq r$. Compared to the finitely-supported signed measures considered in Section 1.3, this introduces additional information that is local to each of the points ξ_1, \dots, ξ_r .

In the course of this chapter, we first examine the moment variety corresponding to first-order local Dirac distributions and relate it to the moment variety of Pareto distributions. Afterwards, we investigate the problem of parameter recovery and formulate a numerical reconstruction algorithm. We finish by illustrating the reconstruction on some concrete examples.

2.1 Moment variety

An interesting geometric object related to a family of distributions is its moment variety. It is the algebraic variety that is described by the set of moment vectors that can arise for any instance of a distribution in the family. As such, it can be defined whenever the parametric descriptions of the moments are algebraic expressions in the parameters. Studying the moment variety allows to deduce properties such as rational or algebraic identifiability, which are important for the parameter recovery problem which we investigate in Section 2.3. In the following, we summarize results from [GW20] about the moment variety of local Dirac mixtures.

Definition 2.1.1. Let $f_\alpha \in \mathbb{k}(x_1, \dots, x_s)$, $\alpha \in A$, be a family of rational functions, for some finite indexing set $A \subseteq \mathbb{N}^n$, $n \in \mathbb{N}$. The Zariski closure of the set parametrized by

$$\{[m_\alpha]_{\alpha \in A} \in \mathbb{P}^{\#A-1} \mid m_\alpha = f_\alpha(x) \text{ for } x \in \mathbb{k}^s, \alpha \in A\} \quad (2.1)$$

is a projective variety of dimension at most s . If the f_α are expressions for the α -th moments of a family of n -variate distributions with s parameters, every point in the parametrically given set is a *moment vector* of an element of the family and the variety is called *moment variety* with respect to A of the family of distributions.

Example 2.1.2. Let $d, n \in \mathbb{N}$ and define the indexing set $A := \{\alpha \in \mathbb{N}^n \mid |\alpha| \leq d\}$. We consider the family of Dirac distributions δ_ξ supported at points $\xi \in \mathbb{k}^n$, which depends on $s := n$ parameters, namely the coordinates of ξ . For each $\xi \in \mathbb{k}^n$, the α -th moment of δ_ξ is given by $\int x^\alpha \delta_\xi(x) = \xi^\alpha$ for $\alpha \in A$. Then, with $f_\alpha := x^\alpha$, the Zariski closure of the set parametrized by (2.1), i. e. the moment variety with respect to A of the family of Dirac distributions, is the same as the d -uple Veronese variety defined in Definition 1.5.2. \diamond

From now on, we primarily focus on univariate distributions, where $n = 1$, with the indexing set $A = \{0, \dots, d\}$ for some $d \in \mathbb{N}$. In this case, the moment variety is parametrized by moments up to degree d and we often denote this variety by \mathcal{X} .

Algebraically, the moment variety can be described by its vanishing ideal, the *moment ideal*, which we view as an ideal in the polynomial ring $\mathbb{k}[M_0, \dots, M_d]$. For the family of first-order local mixtures of Diracs, the following theorem gives polynomial generators of

the corresponding moment ideal, which was first proved in [Eis92] using methods from representation theory. A proof employing algebraic and combinatorial methods is given in [GW20]. The variety can also be described as the tangent variety of the Veronese curve, the one-dimensional Veronese variety.

Theorem 2.1.3 ([Eis92, Section 3], [GW20, Theorem 3.1]). *For $d \geq 6$, let $J_d \subseteq \mathbb{k}[M_0, \dots, M_d]$ be the ideal generated by the $\binom{d-2}{2}$ relations*

$$f_{i,j} := (j - i + 3)M_i M_j - 2(j - i + 2)M_{i-1} M_{j+1} + (j - i + 1)M_{i-2} M_{j+2},$$

for $2 \leq i \leq j \leq d - 2$. Then J_d is equal to the homogeneous ideal of the moment variety of the family of first-order local mixtures of Dirac distributions.

2.2 Pareto distribution

In this section, we establish a connection between the moments of the Pareto distribution and the moments of first-order local mixtures of Diracs. The Pareto distribution, named after Vilfredo Pareto, is a probability distribution that has many useful practical applications; see for instance [Arn83]. In the univariate case, its probability density function is given by

$$\varphi(x) := \frac{\theta \xi^\theta}{x^{\theta+1}} \mathbb{1}_{\{x \geq \xi\}},$$

where $\theta, \xi \in \mathbb{R}_{>0}$. The moments of this distribution are

$$m_i = \begin{cases} \frac{\theta}{\theta-i} \xi^i, & i < \theta, \\ \infty, & i \geq \theta; \end{cases}$$

see [Arn83]. Below, we show that it is possible to choose a different parametrization for the moments of the Pareto distribution in such a way that the moments are the reciprocals of the first-order local mixtures of Diracs. We use this fact to obtain generators of the vanishing ideal of the moment variety corresponding to the family of Pareto distributions.

2.2.1 Ideal generators

Algebraically, we are only interested in the cases in which the moments of the Pareto distribution are finite. These are described by the image of the map

$$\begin{aligned} \mathbb{R}_{>d} \times \mathbb{R}_{>0} &\longrightarrow \mathbb{P}^d, \\ (\theta, \xi) &\longmapsto [m_0 : \dots : m_d] = \left[\frac{\theta}{\theta-i} \xi^i \right]_{0 \leq i \leq d}, \end{aligned}$$

for a given $d \in \mathbb{N}$, where \mathbb{P}^d denotes the projective space over \mathbb{C} . Although the image has real coordinates, we work over the complex numbers, the algebraic closure of \mathbb{R} .

Note that the moments are rational functions in the two parameters θ, ξ . Thus, we can define the moment variety of the Pareto distribution as the Zariski closure over \mathbb{C} of the image of the above map. Since $\mathbb{R}_{>d}$ is Zariski-dense in \mathbb{C} , we may extend the domain of the parametrization to $(\mathbb{C} \setminus \{0, \dots, d\}) \times \mathbb{C}^*$ without changing the Zariski closure of the image. Let ρ denote the extended map $\rho: (\mathbb{C} \setminus \{0, \dots, d\}) \times \mathbb{C}^* \rightarrow \mathbb{P}^d$.

With this notation, the following result establishes the connection between the moment varieties of Pareto distributions and first-order local mixtures of Dirac distributions. Note that two varieties \mathcal{X}, \mathcal{Y} are *birationally equivalent* if there exist two rational maps $\mathcal{X} \dashrightarrow \mathcal{Y}$, $\mathcal{Y} \dashrightarrow \mathcal{X}$ such that both compositions are the identity as rational maps. In this case, there is an isomorphism between Zariski-dense open subsets of \mathcal{X} and \mathcal{Y} . For further details, we refer to [Har87, Section 1.4].

Proposition 2.2.1. *Let $\mathcal{Y} := \overline{\text{im}(\rho)} \subseteq \mathbb{P}^d$ be the Pareto moment variety and $\mathcal{X} \subseteq \mathbb{P}^d$ the moment variety of 1-local mixtures of Diracs, that is, the tangent variety of the Veronese curve. Then, for a suitable choice of parametrization, \mathcal{X} and \mathcal{Y} are birationally equivalent via the rational map*

$$\psi: \mathbb{P}^d \dashrightarrow \mathbb{P}^d, \quad [m_0 : \dots : m_d] \mapsto [m_0^{-1} : \dots : m_d^{-1}].$$

Proof. We change the parametrization of the Pareto moment variety via the bijective map

$$\begin{aligned} \eta: \{(\theta, \xi) \in \mathbb{C}^* \times \mathbb{C}^* \mid -\xi\theta^{-1} \neq 1, \dots, d\} &\longrightarrow (\mathbb{C} \setminus \{0, \dots, d\}) \times \mathbb{C}^*, \\ (\theta, \xi) &\longmapsto (-\xi\theta^{-1}, \xi^{-1}), \end{aligned}$$

which leaves \mathcal{Y} as the closure of the image of $\rho \circ \eta$ unchanged. With this parametrization, the moments are of the form

$$[m_0 : \dots : m_d] = \rho(\eta(\theta, \xi)) = \left[\frac{-\xi\theta^{-1}}{-\xi\theta^{-1} - i} \xi^{-i} \right]_{0 \leq i \leq d} = \left[\frac{1}{\xi^i + i\theta\xi^{i-1}} \right]_{0 \leq i \leq d},$$

so ψ maps points from the image of $\rho \circ \eta$ to moment vectors of 1-local mixtures of Diracs, that is, to points on \mathcal{X} . Then $\psi|_{\mathcal{Y}}: \mathcal{Y} \dashrightarrow \mathcal{X}$ is a rational map that is an isomorphism on the Zariski-dense subset $\text{im}(\rho \circ \eta)$ of \mathcal{Y} , as $\text{im}(\rho \circ \eta) \subseteq \{m_i \neq 0\}$. Being the tangent variety of an irreducible variety, \mathcal{X} is irreducible by [Lan12, Section 8.1]. Thus, the image of $\psi|_{\mathcal{Y}}$ is dense in \mathcal{X} , which proves the claim. \square

Generators for the ideal defining the moment variety of Pareto distributions are given by the following theorem. The variety is contained in \mathbb{P}^d ; we describe the affine subset with $m_0 \neq 0$.

Theorem 2.2.2. *For $d \geq 6$, let $\tilde{I}_d^{\text{inv}} \subseteq R := \mathbb{C}[M_1, \dots, M_d]$ be the ideal generated by the $\binom{d-2}{2}$ polynomials*

$$(j-i+3)M_{i-2}M_{i-1}M_{j+1}M_{j+2} - 2(j-i+2)M_{i-2}M_iM_jM_{j+2} + (j-i+1)M_{i-1}M_iM_jM_{j+1}$$

for $2 \leq i \leq j \leq d-2$, where $M_0 := 1$. Then the affine Pareto moment ideal is equal to the saturation

$$\tilde{I}_d^{\text{inv}} : (M_1 \cdots M_d)^\infty.$$

Proof. Let $\tilde{I}_d \subseteq R = \mathbb{C}[M_1, \dots, M_d]$ be the dehomogenization of the moment ideal of first-order local mixtures of Diracs which was studied in Theorem 2.1.3 (cf. [GW20, Section 3]). In order to restrict to the algebraic torus where the rational map ψ given in Proposition 2.2.1 is defined, consider $J := R[y]\tilde{I}_d + \langle M_1 \cdots M_d y - 1 \rangle \subseteq R[y]$. The restriction of the map ψ to the torus agrees with the torus automorphism induced by the homomorphism

$$\begin{aligned} \bar{\psi}: R[y] &\longrightarrow R[y], \\ y &\longmapsto M_1 \cdots M_d, \\ M_i &\longmapsto M_1 \cdots M_{i-1} M_{i+1} \cdots M_d y, \quad \text{for } 1 \leq i \leq d. \end{aligned}$$

Note that we can choose an ideal $I' \subseteq R$ such that in $R[y]$ we have the equality $\bar{\psi}(J) = R[y]I' + \langle M_1 \cdots M_d y - 1 \rangle$ by observing that, for any $f \in \tilde{I}_d \subseteq R$, we can choose a suitable $k \in \mathbb{N}$ such that

$$(M_1 \cdots M_d)^k \bar{\psi}(f) \equiv f' \pmod{\langle M_1 \cdots M_d y - 1 \rangle}$$

for some $f' \in R$. In particular, this construction establishes a bijection between the generating set of \tilde{I}_d given in Theorem 2.1.3 and the generating set of \tilde{I}_d^{inv} . Therefore, we choose $I' := \tilde{I}_d^{\text{inv}}$. In order to describe the affine closure of the image of ψ , we intersect $\bar{\psi}(J)$ with R , which is equal to

$$\bar{\psi}(J) \cap R = \tilde{I}_d^{\text{inv}} : (M_1 \cdots M_d)^\infty$$

by [CLO15, Theorem 4.4.14] from which we conclude. \square

Note that taking the saturation in the construction is necessary, since otherwise the variety could have additional irreducible components that are supported on the boundary of the algebraic torus, which are not part of the Pareto moment variety.

2.3 Recovery of parameters

While in the previous sections we have focused on local mixtures of a Dirac distribution at a single point, we now consider *mixtures of* local mixtures of Dirac distributions at multiple points and study the problem of parameter recovery from moments. More precisely, we consider distributions of the form

$$\sum_{j=1}^r \lambda_{j0} \delta_{\xi_j} + \lambda_{j1} \delta'_{\xi_j} + \cdots + \lambda_{j,l_j} \delta_{\xi_j}^{(l_j)}$$

for r points $\xi_1, \dots, \xi_r \in \mathbb{k}$ and weights $\lambda_{j,k_j} \in \mathbb{k}$, for $0 \leq k_j \leq l_j$, $1 \leq j \leq r$.

In [GW20, Section 5.2], a symbolic approach for recovering the parameters of a mixture of local Dirac mixtures from its moments is described which is based on elementary symmetric polynomials and relies on elimination theory. This is mainly useful from a theoretical point of view, as only small problems can be solved by this approach in practice, due to the high complexity of the Gröbner basis computations involved in this.

In the following, we describe an algorithm for parameter recovery that is motivated by Prony's method and which can be implemented numerically. Compared to the approach based on elimination theory, this significantly widens the kind of problems that can be solved. This is explained in more detail in Example 2.3.3 as well as Remarks 2.3.4 and 2.3.5. The contents are adapted from [GW20, Section 5.3]. We also refer to [Mou18] for related results; in particular, we closely follow the discussion of Prony's method as it covers the case in which the Prony ideal has multiplicities, which we have not addressed in Section 1.3. The variant of Prony's method that we use here is summed up in Theorem 2.3.1 below.

Let us first fix some notation, following the exposition in Section 1.3.1. Let $R = \mathbb{k}[X]$ be the univariate polynomial ring and denote by $R_{\leq d}$ the vector subspace of polynomials of degree at most $d \in \mathbb{N}$. Further, let $R^* := \text{Hom}(R, \mathbb{k})$ be the dual \mathbb{k} -vector space of the polynomial ring R . Given any sequence $(m_i)_{i \in \mathbb{N}}$, $m_i \in \mathbb{k}$, define $\sigma \in R^*$ to be the \mathbb{k} -linear functional

$$\sigma: R \longrightarrow \mathbb{k}, \quad X^i \longmapsto m_i. \quad (2.2)$$

If the sequence $(m_i)_{i \in \mathbb{N}}$ is the sequence of moments of a distribution, then σ is the moment functional of the distribution. For $d, e \in \mathbb{N}$, denote by $H_{d,e}$ the Hankel matrix with respect to the monomial basis X^i , $0 \leq i \leq e$, and the dual basis $(X^j)^*$, $0 \leq j \leq d$, of the map

$$H_{d,e}: R_{\leq e} \longrightarrow R_{\leq d}^*, \quad p \longmapsto (q \mapsto \sigma(pq)). \quad (2.3)$$

Hence, the matrix is of the form

$$H_{d,e} = (m_{i+j})_{0 \leq i \leq d, 0 \leq j \leq e}$$

and is of size $(d+1) \times (e+1)$. If $(m_i)_{i \in \mathbb{N}}$ is a sequence of moments of some distribution, $H_{d,e}$ is the *moment matrix*.

Assume now we are given an r -mixture of local Dirac mixture distributions located at r points $\xi_1, \dots, \xi_r \in \mathbb{k}$. Then its moments are of the form

$$m_i = \sum_{j=1}^r \sum_{k_j=0}^{l_j} \lambda_{j,k_j} \frac{i!}{(i-k_j)!} \xi_j^{i-k_j} = \sum_{j=1}^r (\Lambda_j(\partial)(X^i))(\xi_j) \in \mathbb{k}, \quad i \in \mathbb{N}, \quad (2.4)$$

where $\lambda_{j,k_j} \in \mathbb{k}$, $0 \leq k_j \leq l_j$, $1 \leq j \leq r$, are suitable parameters and $\Lambda_j(\partial) := \sum_{k_j=0}^{l_j} \lambda_{j,k_j} \partial^{k_j} \in \mathbb{k}[\partial]$ is a polynomial of degree l_j in the variable ∂ that is applied

to the monomial X^i as a differential operator, i. e. $\partial X^i = iX^{i-1}$ for $i \in \mathbb{N}$. We cite the following theorem in order to rephrase it in our language. In particular, we specialize it to the univariate case. We remark that the theorem can also be understood in terms of the canonical form of a binary form. For a detailed treatment of this viewpoint, we refer to [IK99] and the references therein.

Theorem 2.3.1 ([IK99, Theorem 1.43], [Mou18, Theorem 4.1]). *Let \mathbb{k} be an algebraically closed field of characteristic 0, let $R = \mathbb{k}[X]$ and let $m_0, m_1, \dots, m_{2s} \in \mathbb{k}$ for some $s \in \mathbb{N}$. Let $H_{s-1, s-1}, H_{s, s}$ be the corresponding Hankel matrices as in (2.3). Assume $\text{rk } H_{s-1, s-1} = \text{rk } H_{s, s} = r'$. Then there exists a unique mixture of local mixtures of Diracs*

$$\mu := \sum_{j=1}^r \Lambda_j(\partial) \delta_{\xi_j}$$

for some $r \in \mathbb{N}$, $\xi_j \in \mathbb{k}$, $0 \neq \Lambda_j \in \mathbb{k}[\partial]$, such that $\sum_{j=1}^r 1 + \deg(\Lambda_j) = r'$ and its moments up to degree $2s$ coincide with m_0, \dots, m_{2s} . Further, as ideals of R , it holds that

$$\langle \ker H_{s-1, s} \rangle = \bigcap_{j=1}^r \langle X - \xi_j \rangle^{1 + \deg \Lambda_j}.$$

Note that the condition $\sum_{j=1}^r 1 + \deg(\Lambda_j) = r'$ is due to the fact that $\deg(\Lambda_j) + 1$ is the dimension of the vector space spanned by Λ_j and all its derivatives, as we specialized to the univariate case. Moreover, observe that this theorem can be regarded as a variant of Theorems 1.3.1 and 1.3.6. While the latter are statements about radical ideals, Theorem 2.3.1 allows for ideals whose primary components have multiplicities.

As with Prony's method in Algorithm 1.1, Theorem 2.3.1 leads to an algorithm for recovering the parameters of the distribution μ from finitely many moments: first, compute the points ξ_1, \dots, ξ_r from $\ker H_{s-1, s}$ for s sufficiently large; next, determine the weights Λ_j from the linear system described by (2.4). If $\deg \Lambda_j = 0$ for all $j = 1, \dots, r$, this algorithm agrees with the classic Prony method, also referred to as *Sylvester's algorithm* in the classical algebraic geometry literature, which was discussed in Section 1.3. In the context of this chapter, the more interesting setting is thus the case in which $\deg \Lambda_j \neq 0$ for some or all $j = 1, \dots, r$.

In the following, we refine this algorithm for the case of mixtures of local mixtures of Diracs of fixed order $l := l_1 = \dots = l_r$ where $l_j = \deg \Lambda_j$ for $j = 1, \dots, r$. In this case, it is usually possible to recover the parameters from significantly fewer moments. (More generally, when l_1, \dots, l_r are not equal, one can set $l := \max\{l_1, \dots, l_r\}$, but the potential of saving moments is not as high, in this case.)

Theorem 2.3.2. *Let \mathbb{k} be a field of characteristic 0 and let $\mu := \sum_{j=1}^r \Lambda_j(\partial) \delta_{\xi_j}$ be an r -mixture of l -th-order local mixtures of Diracs, i. e. $\xi_j \in \mathbb{k}$ and $\Lambda_j \in \mathbb{k}[\partial]$ with $\deg(\Lambda_j) = l$, $1 \leq j \leq r$. Then, the parameters Λ_j, ξ_j of μ can be recovered from moments $m_0, m_1, \dots, m_{2(l+1)r-1}$ of μ using Algorithm 2.1.*

Proof. Let σ be the functional associated to μ , as defined above in (2.2). Then, it follows from [Mou18, Theorem 3.1, Proposition 3.9] that $\text{rk } H_{d,e} = (l+1)r$ for all $d, e \geq (l+1)r-1$. In particular, for $s := (l+1)r$, we have

$$\text{rk } H_{s-1,s-1} = \text{rk } H_{s,s} = (l+1)r.$$

The algorithm is based on the following observation. Let $p \in R$ be the polynomial $p := \prod_{j=1}^r (X - \xi_j) = X^r + \sum_{i=0}^{r-1} p_i X^i$, noting that knowledge of p is enough to recover the points ξ_j . By the addendum of Theorem 2.3.1, we have $p^{l+1} \in \langle \ker H_{s-1,s} \rangle \otimes_{\mathbb{k}} \bar{\mathbb{k}}$ where $\bar{\mathbb{k}}$ is the algebraic closure of \mathbb{k} . Since also $p^{l+1} \in R$, it follows in particular that p is mapped to 0 under the composition of the maps

$$\begin{array}{ccccc} R_{\leq r} & \longrightarrow & R_{\leq (l+1)r} & \longrightarrow & R_{\leq r-1}^* \\ q & \longmapsto & q^{l+1} & \longmapsto & H_{r-1,(l+1)r} q^{l+1}, \end{array}$$

where the second map is the \mathbb{k} -linear map given by the moment matrix $H_{r-1,(l+1)r}$, which is a submatrix of $H_{s-1,s}$. The first map however is non-linear, defined by taking the $(l+1)$ -th power of q viewed as a polynomial.

For the polynomial p , this yields the following polynomial system of r equations of degree $l+1$ in r variables p_0, \dots, p_{r-1} which are the monomial coefficients of p :

$$H_{r-1,(l+1)r} p^{l+1} = 0. \tag{2.5}$$

Note that p is monic. By Bézout's theorem, this system of equations either has infinitely many or at most $(l+1)^r$ solutions. If the solution set is infinite, we need to add more algebraic constraints to the system in order to determine p , which is done by adding more rows to the moment matrix.

By hypothesis, we have $\xi_1, \dots, \xi_r \in \mathbb{k}$. Therefore, termination of this algorithm and correct recovery of the points ξ_1, \dots, ξ_r follow from Theorem 2.3.1.

As for computation of the weights λ_{jk} in Step 5, note that, once the roots ξ_j have been computed, the moments are a linear combination of the monomials ξ_j^i and their derivatives given by Equation (2.4), so to compute the weights λ_{jk} , we solve the linear system

$$(V_1, \dots, V_r) \begin{pmatrix} \lambda_{10} \\ \vdots \\ \lambda_{1l} \\ \vdots \\ \lambda_{rl} \end{pmatrix} = \begin{pmatrix} m_0 \\ \vdots \\ m_d \end{pmatrix}$$

for $d \geq s$, where (V_1, \dots, V_r) is a *confluent Vandermonde matrix*, for which each block is

given by

$$V_j = \left((\partial^k X^i)(\xi_j) \right)_{\substack{0 \leq i \leq d \\ 0 \leq k \leq l}} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ \xi_j & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \xi_j^d & d\xi_j^{d-1} & \cdots & \frac{d!}{(d-l)!}\xi_j^{d-l} \end{pmatrix}.$$

Since the system is linear, uniqueness of the solution follows from the claim that the confluent Vandermonde matrix is of full rank s . Without loss of generality, we can assume that the confluent Vandermonde matrix is of size $s \times s$ by choosing a suitable submatrix. Then the claim follows from the fact that the Hermite interpolation problem has a unique solution if the points ξ_1, \dots, ξ_r are distinct or, equivalently, from the product formula for the determinant of a square confluent Vandermonde matrix; see [HJ94, Problem 6.1.12]. \square

Algorithm 2.1 Parameter recovery for mixtures of local Dirac mixtures with r components of order l

Input: The (maximum) order $l \geq 0$ of the mixture components, the number of mixture components $r \geq 1$ and the moments $m_0, \dots, m_{(l+2)r}, \dots$

Output: The parameters ξ_j and λ_{jk} for $1 \leq j \leq r$, $0 \leq k \leq l$, satisfying (2.4).

- 1: Solve the polynomial system $H_{r-1, (l+1)r} p^{l+1} = 0$ for p .
 - 2: If the solution set is infinite, increase the number of rows of the moment matrix and repeat.
 - 3: If there is more than one solution, use further information, such as the additional moment $m_{(l+2)r}$, to restrict to a single solution p .
 - 4: Compute the roots ξ_1, \dots, ξ_r of p .
 - 5: Compute the weights λ_{jk} by solving a confluent Vandermonde system.
-

Note that the algorithm is designed to use as few moments as possible. See also Remark 2.3.5 for a discussion of the number of moments used by this algorithm.

Example 2.3.3. For $r = 2$, $l = 1$, let m_0, \dots, m_5 be the moments of a corresponding distribution and write $p = p_0 + p_1 X + X^2$. Then the system of equations (2.5) is given by the quadratic equations

$$\begin{pmatrix} m_0 & m_1 & m_2 & m_3 & m_4 \\ m_1 & m_2 & m_3 & m_4 & m_5 \end{pmatrix} \begin{pmatrix} p_0^2 \\ 2p_0 p_1 \\ 2p_0 + p_1^2 \\ 2p_1 \\ 1 \end{pmatrix} = 0. \quad (2.6)$$

If ξ_1, ξ_2 are the points of the underlying distribution, one solution of this system is given by $p = (X - \xi_1)(X - \xi_2)$, that is, $p_1 = -(\xi_1 + \xi_2)$ and $p_0 = \xi_1 \xi_2$. These are elementary symmetric polynomials in ξ_1, ξ_2 . Hence, computing p_1 by eliminating p_0 from the system (2.6), and vice versa, is equivalent to the process of recovering the parameters from elementary symmetric polynomials detailed in [GW20, Section 5.2]. However, with

the approach presented here, we need to eliminate only a single variable instead of 5, which makes this much more viable from a computational point of view. For larger problems, solving the non-linear polynomial system in Step 1 of Algorithm 2.1 is still a very challenging problem, though. \diamond

Remark 2.3.4. An advantage of the approach based on elimination theory in [GW20, Section 5.2] is that it works with symbolic moments. Thus, for a particular choice of r and l , the result gives a solution for any choice of moments. In contrast, the approach based on Prony’s method presented here only computes a solution for a fixed set of moments, so that only one particular instance of a parameter recovery problem is solved, but has the advantage of being computationally more tractable. Depending on the application, either approach may be useful. \diamond

Remark 2.3.5. We discuss some of the steps involved in Algorithm 2.1 in more detail. Solving the system in Step 1 can be done using symbolic algebraic methods, which usually involve the computation of a Gröbner basis, or using numerical tools. In Example 2.4.2 for instance, we use a numerical solver for this which is based on homotopy continuation methods. See for instance [Li03] for an introduction to homotopy continuation.

Restricting from finitely many solutions to a single one using the additional moment $m_{(l+2)r}$ in Step 3 works by observing that p is also a solution to the larger system $H_{r,(l+1)r}p^{l+1} = 0$. If a numerical solver is used, the computed solution will only be approximately zero, and one should assert that the selected solution is significantly closer to zero than all other possible choices. Another common approach to check uniqueness of the solution is to perform monodromy loop computations using a homotopy solver.

An upper bound for the number of moments used by the algorithm is $2(l+1)r$, since the moments $m_0, \dots, m_{2(l+1)r-1}$ are always enough for recovery, as stated in Theorem 2.3.2. Then solving the polynomial system in Algorithm 2.1 simplifies, since the solution is in the kernel of $H_{(l+1)r-1,(l+1)r}$ which is a linear problem. In this case, the algorithm performs the same computation as [Mou18, Algorithm 3.2], so this guarantees termination.

However, as Algorithm 2.1 solves a more specific problem, it can usually recover the parameters using a much smaller number of moments. The polynomial system in Step 1 consists of r equations of degree $l+1$ in r unknowns, so, generically, we expect finitely many solutions in Step 2 already in the first iteration of the algorithm. This means we expect to algebraically identify the parameters from the moments $m_0, \dots, m_{(l+2)r-1}$ and to rationally identify them using one additional moment, so usually we do not need all the moments up to $m_{2(l+1)r-1}$. This is also what we observe in practice, so we do not seem to get infinitely many solutions for generic input if we use the moments up to $m_{(l+2)r-1}$. By a parameter count, we cannot expect to recover the parameters from fewer moments, so the number of moments we use in practice is the minimal number possible. \diamond

Here the term *algebraic identifiability* means that the map from the parameters to the moments is generically finite-to-one; see for instance [ARS18]. Likewise, *rational identifiability* holds if the moment map is generically one-to-one. Since, by Theorem 2.3.2,

it is possible to uniquely recover the parameters from the moments $m_0, \dots, m_{2(l+1)r-1}$, rational identifiability certainly holds if $d \geq 2(l+1)r - 1$, but usually one expects this to be the case already for much smaller d . The following proposition states that algebraic identifiability holds as soon as $d \geq (l+2)r - 1$, so one expects rational identifiability to hold if $d \geq (l+2)r$, as explained in Remark 2.3.5.

Proposition 2.3.6. *Let $d \geq (l+2)r - 1$. Then algebraic identifiability holds for the moment map sending the parameters ξ_j, Λ_j to the moments m_0, \dots, m_d , where $\deg \Lambda_j = l$, $1 \leq j \leq r$.*

Proof. By [CGG02, Proposition 3.1], the secant varieties of the tangent variety of the Veronese curve are non-defective, that is, for $l = 1$, the dimension of the moment variety in \mathbb{P}^d for mixtures with r components of order l is the expected one: $\min(3r - 1, d)$. In particular this means that the moment variety fills the ambient space sharply if $d = 3r - 1$ and does not fill the ambient space if $d > 3r - 1$. Thus, the moment map is generically finite-to-one if $d \geq 3r - 1$. Note that for $d < 3r - 1$ the cardinality of the preimage of a generic point is infinite for dimension reasons.

Similarly, for $l \geq 2$, the moment variety is a secant variety of the l -th osculating variety to the Veronese curve which is non-defective by [BCGI07, Section 4], so the moment map is generically finite-to-one for $d \geq (l+2)r - 1$. \square

Remark 2.3.7. It would be interesting to generalize Algorithm 2.1 to the multivariate setting. Note that in this case [Mou18, Algorithm 3.2] can be used to find the decomposition. However, since this does not take into account the special structure of our input, namely that all the mixture distributions have the same order l , this approach might use more moments than necessary. This is similar to the univariate case as explained in Remark 2.3.5.

Further, the algorithm in [BT20, Section 6] also computes a generalized decomposition from a given set of moments. This algorithm differs from our Algorithm 2.1 in that it computes parameters of *any* generalized decomposition explaining the given moments, rather than the unique decomposition in which each term corresponds to the same order l . In the one-dimensional case, when using as few moments as possible, this usually leads to a non-generalized decomposition, which does not recover the parameters we are interested in. See also the related discussion in [BT20, Section 7.1]. \diamond

Remark 2.3.8. We briefly discuss how the problem of parameter recovery of a mixture of 1-local mixtures of Diracs simplifies, if the mixture components $\delta_{\xi_j} + \theta_j \delta_{\xi_j}^l$, $1 \leq j \leq r$, are known to differ only in the parameters ξ_j , but have a constant parameter $\theta := \theta_1 = \dots = \theta_r$. For this, it is convenient to denote the moments of a distribution X by $\mathbb{E}(X^i)$ for $i \in \mathbb{N}$, as it underlines the algebraic relationship between the moments of the two distributions we consider here.

Let us assume that X is a distribution which has moments that are of the form $\mathbb{E}(X^i) = \sum_{j=1}^r \lambda_j (\xi_j^i + \theta i \xi_j^{i-1})$ with a fixed parameter θ . Further, let Y be the distribution with moments $\mathbb{E}(Y^i) = \sum_{j=1}^r \lambda_j \xi_j^i$, so Y is a linear combination of Dirac distributions. Then

we have

$$\mathbb{E}(X^i) = \mathbb{E}(Y^i) + \theta i \mathbb{E}(Y^{i-1})$$

and conversely

$$\mathbb{E}(Y^i) = \sum_{k=0}^i \frac{i!}{k!} (-\theta)^{i-k} \mathbb{E}(X^k).$$

Hence, if θ is known, this allows to recover the moments of the distribution Y from the moments of X . The parameters of Y can then be recovered, e. g. using Prony's method.

In case θ is fixed, but unknown, treating θ as a variable in the moment matrix $H_r(Y) := (\mathbb{E}(Y^{i+j}))_{0 \leq i, j \leq r}$, it can be determined as one of the roots of $\det H_r(Y)$, which is a polynomial of degree $r(r+1)$ in θ . \diamond

2.4 Applications

2.4.1 Moments and Fourier coefficients

In this section, we show how the tools developed in this chapter can be applied to the problem of recovering a piecewise-polynomial function supported on the interval $[-\pi, \pi]$ from Fourier coefficients; see [PT14]. For this, we describe how moments of a mixture of local mixtures of Diracs arise as the Fourier coefficients of a piecewise-polynomial function and illustrate this numerically. For simplicity, we focus on the case $l = 1$ of piecewise-(affine-)linear functions.

Let $t_j \in [-\pi, \pi]$, $1 \leq j \leq r$, be real points and let $f: [-\pi, \pi] \rightarrow \mathbb{C}$ be the piecewise-linear function given by

$$f(x) := \sum_{j=1}^{r-1} (f_j + (x - t_j) f'_j) \mathbb{1}_{[t_j, t_{j+1})}(x), \quad (2.7)$$

where $f_j, f'_j \in \mathbb{C}$. In particular, splines of degree 1 are of this form, but we do not require continuity here. The Fourier coefficients of f are defined to be

$$c_k := \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) e^{-ikx} dx, \quad k \in \mathbb{Z},$$

from which we obtain

$$c_k = \frac{1}{2\pi(i k)^2} \sum_{j=1}^r (i k (f_j - f_{j-1} + (t_{j-1} - t_j) f'_{j-1}) + (f'_j - f'_{j-1})) e^{-i k t_j},$$

for $k \in \mathbb{Z} \setminus \{0\}$, where $f_0, f'_0, f_r, f'_r := 0$. Further, let

$$\begin{aligned} \xi_j &:= e^{-i t_j}, \\ \lambda_j &:= \xi_j^{-s} (f'_j - f'_{j-1} - i s (f_j - f_{j-1} + (t_{j-1} - t_j) f'_{j-1})), \\ \lambda'_j &:= \xi_j^{1-s} i (f_j - f_{j-1} + (t_{j-1} - t_j) f'_{j-1}). \end{aligned} \quad (2.8)$$

Assume now, we are given finitely many Fourier coefficients c_{-s}, \dots, c_s for some $s \in \mathbb{N}$. Then, for $0 \leq k \leq 2s$, $k \neq s$, we define

$$m_k := 2\pi(i(k-s))^2 c_{k-s} = \sum_{j=1}^r \lambda_j \xi_j^k + \lambda'_j k \xi_j^{k-1}. \quad (2.9)$$

Thus, from the knowledge of Fourier coefficients c_{-s}, \dots, c_s of f , we can compute m_k , $k \neq s$, which we interpret as the moments of a mixture of 1-local mixtures of Diracs with support points ξ_j on the unit circle. Extending the definition to m_s , by construction we have

$$m_s := \sum_{j=1}^r \lambda_j \xi_j^s + \lambda'_j s \xi_j^{s-1} = \sum_{j=1}^r f'_j - f'_{j-1} = 0.$$

All in all, we know the moments m_0, \dots, m_{2s} of this 1-local mixture. Recovering the parameters $\xi_j, \lambda_j, \lambda'_j$ via Algorithm 2.1 generically requires the moments m_0, \dots, m_{3r} , so we need to choose $2s \geq 3r$. Subsequently retrieving the original parameters t_j, f_j, f'_j from (2.8) is straightforward.

Remark 2.4.1. The piecewise-linear function f is viewed as a periodic function on the interval $[-\pi, \pi)$, in the discussion above. For simplicity in presentation, we assumed that f is constantly zero outside of $[t_1, t_r)$, representing a constant line segment. More generally, one can adapt the computation to account for an additional (non-zero) line segment there, without changing the number of jumping points or required samples. Thus, the function f consists of r line segments and has r discontinuities. \diamond

Example 2.4.2. We apply the process described above to a piecewise-linear function with $r = 10$ line segments on the interval $[-\pi, \pi)$. The parameters t_j, f_j, f'_j defining the function as in (2.7) are listed in Table 2.1. The random jump points t_j are chosen uniformly on the interval. The jump heights f_j and slopes f'_j are chosen with respect to a Gaussian distribution. The function as well as the sampling data are visualized in Figure 2.1. Our samples consist of evaluations of the Fourier partial sum of the function at equidistantly-spaced sampling points. By Fourier transform, these carry the same information as the Fourier coefficients of the piecewise-linear function (cf. [PPST18, Section 3.1.3]). The number of sampling points equals the number of Fourier coefficients needed for reconstruction, namely $3r + 1 = 31$. From $2s \geq 3r$, we determine $s = 15$. We compute the Fourier coefficients c_{-s}, \dots, c_s from the given data and add some noise to each of these coefficients, sampled from a Gaussian distribution with standard deviation 10^{-12} .

In order to reconstruct the piecewise-linear function from the Fourier coefficients, we compute the moments m_0, \dots, m_{3r} via Equation (2.9) and apply Algorithm 2.1 using numerical tools. From the moments m_0, \dots, m_{3r-1} , we get a system of r quadratic equations in r unknowns, which we solve using the Julia package *HomotopyContinuation.jl* [BT18], version 0.3.2, from which we obtain up to 2^r finite solutions. From these, we choose the one that best solves the quadratic equation system $H_{r,2r} p^2 = 0$ induced by the additional moment m_{3r} . In this example, we obtain 1024 solutions, the best of which




Figure 2.1: The piecewise-linear function of Example 2.4.2 with $r = 10$ line segments (solid); the Fourier partial sum approximation of order $s = 15$ (dotted) and $2s + 1 = 31$ equidistantly-spaced sampling points.

has error $1.54 \cdot 10^{-10}$ in the ℓ^2 -norm; the second best solution has error $3.70 \cdot 10^{-4}$, which is significantly larger, so we accept the solution.

Next, we compute the points ξ_j using the Julia package *PolynomialRoots.jl* [SG12], version 0.2.0, and solve an overdetermined confluent Vandermonde system for the weights λ_j, λ'_j , for which we use a built-in least-squares solver. Lastly, we use (2.8) to compute the parameters t_j, f_j, f'_j . Julia code for these computations is included in [Wag21]. The numerical computations were carried out using the Julia language [BEKS17], version 1.0.0.

In this example, the total error we get for the reconstructed points t_1, \dots, t_{10} is $3.89 \cdot 10^{-10}$ in the ℓ^2 -norm, whereas for f_j and f'_j , $1 \leq j \leq 9$, we get $2.15 \cdot 10^{-7}$ and $2.35 \cdot 10^{-7}$, respectively. Even though, in this example, one of the line segments is quite far off of

Table 2.1: The parameters of the piecewise-linear function of Example 2.4.2.

j	t_j	f_j	f'_j
1	-2.814030328751694	-0.20121264876344414	-0.775069863870378
2	-2.457537611167516	-0.35221920435611676	-0.9795392068942285
3	-1.4536804635810938	-0.9254256123988903	0.26040229778962753
4	-1.1734228328971805	0.4482105605664995	-0.46848914917290574
5	-0.6568874684874002	1.11978779941218	-0.8808972481620518
6	0.54049294753688	0.3012272070859375	0.2777255506414151
7	1.0213620344785337	-0.8357295816882367	1.5239161501048377
8	1.0930147137662223	-0.2071744440917742	-1.7777276640658903
9	1.6867064885416054	0.8681006042361324	-2.9330595087256466
10	2.7678373800858678		

the Fourier partial sum, as shown in Figure 2.1, the sampling data still contains enough information to reconstruct it. \diamond

We observe that we cannot always reconstruct the randomly chosen points correctly using homotopy continuation, but many times reconstruction is successful. We expect that the separation distance among the points plays a major part in numerical reconstruction. If the randomly chosen points are badly separated, it will be difficult to distinguish them numerically by just using the moments, as is the case if $l = 0$; see [KN20; BDGY21].

Further, we observe that, after having obtained the points, solving the confluent Vandermonde system often induces additional errors of about three orders of magnitude, resulting from the possibly bad condition of the confluent Vandermonde matrix. A detailed discussion of this condition number exceeds the scope of this work, so we leave it for further study.

Remark 2.4.3. If $l > 1$, one can adapt Algorithm 2.1 in a similar fashion to the reconstruction of functions that are piecewise defined by polynomials of degree l . As this requires solving a system of polynomial equations of degree l , the involved computations are more challenging. Note however that, under additional assumptions on the smoothness of the function, computations can be reduced to a polynomial system of smaller degree. For example, if we let $l = 3$ and additionally impose C^1 -continuity, the second derivative is piecewise-linear, so reconstruction can be accomplished by applying the method outlined above. \diamond

2.4.2 Local mixture distributions

In statistical applications, local mixture models can be used to account for small variations in the data that otherwise are not directly reflected in a particular statistical model. This means the original model is enriched by a truncated Taylor expansion of the probability density function. For details, we refer to [AM07] and the references therein as well as to [Gro21, Section 2.5].

Definition 2.4.4. For a regular exponential family $\phi_\xi(x)$, its *local mixture model* is defined as

$$\psi_\xi(x) := \phi_\xi(x) + \sum_{i=1}^l \theta_i \phi_\xi^{(i)}(x),$$

for parameters $\theta_1, \dots, \theta_l$ such that $\psi_\xi(x) \geq 0$ for all x . Here $\phi_\xi^{(i)}$, $1 \leq i \leq l$, denote the i -th derivatives of ϕ_ξ , which are usually not non-negative.

Since the distribution ψ_ξ is non-negative, it is also a measure by [Sch73, Chapter 1.4, Theorem 5]. The local mixture model is a convolution between a local mixture of Dirac distributions and the member of the exponential family that is centered at 0, by [Sch73, Chapter 6.3, Theorem 8]. Thus, for the distribution $\mu := \delta_\xi + \sum_{i=1}^l \theta_i \delta_\xi^{(i)}$ we have

$$\psi_\xi = \phi_0 * \mu. \tag{2.10}$$

Here, the *convolution* $\phi_0 * \mu$ is defined as a distribution satisfying

$$\int_{\mathbb{R}} f d(\phi_0 * \mu) = \int_{\mathbb{R}} \int_{\mathbb{R}} f(x+y)\phi_0(y)dyd\mu(x),$$

for every compactly-supported function $f \in C^l(\mathbb{R})$ (cf. [Sch73, Chapter 6.2]).

Remark 2.4.5. Let $M_1(t), M_2(t)$ be the moment generating functions of the two distributions in the convolution, respectively. Hence, in particular, $M_2(t) := \sum_{k \geq 0} m_k \frac{t^k}{k!}$ where

$$m_k = \xi^k + \sum_{i=1}^{\min\{l,k\}} (-1)^i \theta_i \frac{k!}{k-i!} \xi^{k-i}$$

denote the moments of the local mixture of Diracs. Note that the sign changes are due to the property of the derivative of the Dirac distribution that

$$\int \phi(x) d\delta_{\xi}^{(i)}(x) = (-1)^i \phi^{(i)}(\xi).$$

Then the product $M_1(t)M_2(t)$ is the moment generating function corresponding to the convolution ψ_{ξ} , so the moments m_k of the underlying local Dirac mixture can be computed if the moments of a probability distribution with density ψ_{ξ} as well as the moments corresponding to ϕ_0 are known. Thus, we can reduce the problem of parameter inference to the problem of parameter recovery of a local Dirac mixture. \diamond

A numerical example of a local mixture of Gaussians is given in [Gro21, Section 2.5]. In the following more involved example from [GW20], we numerically apply the process outlined above to Gaussian distributions by considering a mixture of two local mixtures of Gaussians of order 2. The reconstruction is performed using Algorithm 2.1.

Example 2.4.6 (A mixture of two local Gaussians with known common variance). Let ϕ_0 be the density of a standard Gaussian distribution, let ϕ_{ξ_j} be the density of Gaussians centered at the points ξ_j , $j = 1, 2$, and let

$$\psi := \lambda\psi_{\xi_1} + (1 - \lambda)\psi_{\xi_2}$$

be a 2-mixture of local distributions of order 2, where $\psi_{\xi_j} = \phi_{\xi_j} + \theta_{j1}\phi'_{\xi_j} + \theta_{j2}\phi''_{\xi_j}$. The functions $\phi'_{\xi_j}, \phi''_{\xi_j}$ are also called *Hermite functions* (cf. [PPST18, Section 2.2]). Recall that ψ_{ξ_j} can be expressed by a convolution of ϕ_0 and a second-order local mixture of Dirac distributions by (2.10).

For this example, we choose the parameters as follows:

$$\begin{array}{llll} \xi_1 = -1, & \theta_{11} = 0.1, & \theta_{12} = 0.4, & \lambda = 0.6, \\ \xi_2 = 2, & \theta_{21} = -0.2, & \theta_{22} = 0.6. & \end{array}$$

With this choice of θ_{ji} , $1 \leq i, j \leq 2$, the functions ψ_{ξ_j} are non-negative, so they are indeed probability density functions; see [Mar02, Example 4]. We create a sample of size

20,000 from this probability distribution using Mathematica [Wol18] and compute the empirical moments of that sample. Using that, we derive the empirical moments of the underlying 2-mixture of local Dirac distributions as explained in Remark 2.4.5 and then apply Algorithm 2.1 to infer the parameters. We obtain the following values:

$$\begin{aligned} \xi_1 &= -0.98121, & \theta_{11} &= 0.14076, & \theta_{12} &= 0.39268, & \lambda &= 0.59457, \\ \xi_2 &= 1.95600, & \theta_{21} &= -0.20486, & \theta_{22} &= 0.62641. \end{aligned}$$

These are close to the original parameters – the numerical errors are expected due to the fact that the empirical moments only give an approximation of the true underlying moments. Increasing the sample size can result in a better approximation.

Note that for this process the distribution ϕ_0 is assumed to be known. In particular, we need to know its standard deviation or have a way of estimating it in order to relate the moments of the local mixture of Gaussians and Diracs as in Remark 2.4.5. The original distribution and the reconstructed one are shown in Figure 2.2. \diamond




Figure 2.2: The 2-mixture $\psi = \lambda\psi_{\xi_1} + (1 - \lambda)\psi_{\xi_2}$ of local Gaussian distributions of Example 2.4.6 (solid) and its two components $\lambda\psi_{\xi_1}, (1 - \lambda)\psi_{\xi_2}$ (dashed) on the left, as well as the reconstructions on the right, including the original distribution for comparison (dotted).

3 Moment problems on positive-dimensional varieties

In this chapter, we consider a measure that is supported on an algebraic variety of positive (or, more generally, arbitrary) dimension and study the problem of recovering from its moments the defining data of the measure. As this problem is too broad to be solved in general, we focus on specific subproblems, such as recovering the algebraic variety the measure is supported on (Section 3.4) as well as finding approximations of the measure using only finitely many moments (Section 3.5). This constitutes a generalization of the problem considered in Section 1.3, as Prony's method can be viewed as a tool for reconstructing a signed measure supported on a zero-dimensional algebraic variety from finitely many moments.

Our primary interest pertains to measures with support contained in the n -dimensional affine space or in the complex torus \mathbb{T}^n . To every measure μ that has finite moments, we can associate its moment functional σ , a \mathbb{k} -linear map from the polynomial ring (or Laurent polynomial ring) L to the underlying field \mathbb{k} . More abstractly, we can start with any functional $\sigma: L \rightarrow \mathbb{k}$, without referring to moments of a measure. This is the approach we follow in the first sections of the chapter, as it allows us to make purely algebraic statements without need of any measure theoretic arguments. More specifically, if the measure μ is supported on an algebraic variety $V(\mathfrak{a})$ defined by an ideal $\mathfrak{a} \subseteq L$, then the moment functional σ factors via L/\mathfrak{a} . Thus, we consider functionals $\sigma: L \rightarrow \mathbb{k}$ for which $\mathfrak{a} \subseteq \ker \sigma$ holds. The variety defined by the ideal \mathfrak{a} can be positive-dimensional, but this also includes zero-dimensional ideals as a special case.

In Section 3.1, we start with a short general treatment of sesquilinear forms and sesquilinear maps that can be associated to a functional σ . The concept of sesquilinearity is useful in this context as it allows us to treat both cases, that of measures in affine space and on the torus, simultaneously.

We then continue in Section 3.2 by transferring, to the more general setting of this chapter, the Vandermonde factorization of Lemma 1.3.3 that is such an essential ingredient for Prony's method. Section 3.3 consists of a short discussion of aspects that are relevant for the complex torus. In particular, this relates Toeplitz and Hankel moment matrices.

Section 3.4 addresses one of our leading questions, that of recovering the algebraic variety the measure μ is supported on. This can be achieved by using finitely many moments, both for non-negative as well as signed measures, in affine space and on the complex

torus. In case of non-negative measures, the moment matrices are positive-semidefinite which allows for somewhat stronger statements; we illustrate this difference in some examples.

Finally, Section 3.5 deals with measures on the torus, exclusively. We construct several functions from a finite number of moments that reflect different aspects of the original measure. For these functions, we prove several results about qualitative convergence as the number of moments is increased.

Unless explicitly stated otherwise, the measures we consider are assumed to be non-negative and are defined in terms of the Borel σ -algebra.

3.1 Sesquilinearity

In this section, we set up a framework that allows us to treat in a unified way the two different settings of moment problems we are primarily interested in, namely moment problems on affine space and on the torus. See [Sch17, Chapter 2] for a similar approach to these concepts.

Definition 3.1.1. Let R be a ring with a map $-^\circ: R \rightarrow R$ satisfying

- $(x + y)^\circ = x^\circ + y^\circ$,
- $(xy)^\circ = y^\circ x^\circ$,
- $1^\circ = 1$,
- $(x^\circ)^\circ = x$

for all $x, y \in R$. Then the map $-^\circ$ is called *involution* and R is an *involution ring* (also called **-ring*).

An involutive ring A with involution $-^{\circ A}$ that is also an (associative) algebra over a commutative involutive ring R is an *involution algebra* (also called **-algebra*), if the involution satisfies

$$(ra)^{\circ A} = r^\circ a^{\circ A}$$

for all $r \in R$ and $a \in A$. As this property means that there is no ambiguity, we denote the involution on A by $-^\circ$ as well.

A map $f: A \rightarrow A$ is $^\circ$ -*semilinear* if $f(a + b) = f(a) + f(b)$ and $f(ra) = r^\circ f(a)$ holds for all $r \in R$ and $a, b \in A$. The same definition also applies to maps between submodules of A .

A common example for an involutive ring is the field of complex numbers \mathbb{C} with complex conjugation as involution. It is an involutive algebra over \mathbb{R} , where \mathbb{R} is endowed with the trivial involution. More generally, quadratic field extensions give rise to non-trivial involutive algebras over the original field. An example of a non-commutative involutive

3.1 SESQUILINEARITY

algebra over \mathbb{C} consists of square complex matrices of a given size, for which an involution is defined by taking the conjugate transpose of a matrix. Another important example for our discussion is given in Example 3.1.5 below. Also note that any commutative ring (algebra) is an involutive ring (algebra) with respect to the trivial involution which leaves every element unchanged.

Definition 3.1.2. Let \mathbb{k} be a field and A an (associative) algebra over \mathbb{k} . If $F_d \subseteq A$, $d \in \mathbb{N}$, is a family of \mathbb{k} -vector subspaces satisfying

- $F_d \subseteq F_e$ for $d, e \in \mathbb{N}$ with $d \leq e$,
- $A = \bigcup_{d \in \mathbb{N}} F_d$,
- $1 \in F_0$,
- $F_d \cdot F_e \subseteq F_{d+e}$ for $d, e \in \mathbb{N}$,

then A is a *filtered algebra* over \mathbb{k} and the family $\{F_d\}_{d \in \mathbb{N}}$ is called *filtration* of A . For simplicity of notation, we often denote the filtered components of the filtration by $A_{\leq d} := F_d$.

More generally, one can define filtered algebras over commutative rings in the same way, but here we only make use of this notion for algebras over fields. We only work with algebras and filtered algebras that are commutative. More specifically, we only consider filtrations of the multivariate polynomial ring or Laurent polynomial ring over a field.

Example 3.1.3. Let \mathbb{k} be a field and $R = \mathbb{k}[x_1, \dots, x_n]$ be the polynomial ring in n variables over \mathbb{k} , for some $n \in \mathbb{N}$. Then the total degree of polynomials gives rise to a filtration of R where

$$R_{\leq d} = \{p \in R \mid \deg(p) \leq d\}$$

for $d \in \mathbb{N}$. Similarly, we can define a filtration $\{F_d\}_{d \in \mathbb{N}}$, on R in terms of max-degree by

$$F_d = \bigoplus_{\alpha \in \mathbb{N}^n, |\alpha|_{\infty} \leq d} \mathbb{k}x^{\alpha}.$$

Note that all the filtered components of these two filtrations happen to be \mathbb{k} -vector spaces of finite dimension, which is a useful property when it comes to computations.

Now let $\mathfrak{a} \subseteq R$ be an ideal with $1 \notin \mathfrak{a}$ and define $S = R/\mathfrak{a}$. If $\{F_d\}_{d \in \mathbb{N}}$, is any filtration of R , then

$$G_d := F_d/(\mathfrak{a} \cap F_d)$$

is a filtration of S . For this, observe that G_d can be embedded in G_{d+1} , for all $d \in \mathbb{N}$, via the injective map $p + \mathfrak{a} \cap F_d \mapsto p + \mathfrak{a} \cap F_{d+1}$, where $p \in F_d$. \diamond

For the remainder of this section, we assume, for simplicity, that \mathbb{k} is a field of characteristic 0 together with an involution $-^\circ$ that endows \mathbb{k} with the structure of an involutive ring. Moreover, we denote by $R = \mathbb{k}[x_1, \dots, x_n]$ the polynomial ring in finitely many

variables and fix a filtration $\{R_{\leq d}\}_{d \in \mathbb{N}}$ that turns R into a filtered algebra over \mathbb{k} and has the property that $R_{\leq d}$ is a finite-dimensional \mathbb{k} -vector space for every $d \in \mathbb{N}$. Additionally, we assume that $R \subseteq L$ is a \mathbb{k} -subalgebra of an involutive commutative algebra L over \mathbb{k} . The involution on L is denoted by $-\circ$ as well. Typical examples are the following:

Example 3.1.4. If \mathbb{k} is any field, let $L = R$ and define the involutions on \mathbb{k} and L to act trivially. The filtration on R is defined by total degree as

$$R_{\leq d} = \{p \in R \mid \deg(p) \leq d\}$$

for $d \in \mathbb{N}$. Of particular interest is the case when \mathbb{k} is the field of real numbers \mathbb{R} (or a subfield thereof). \diamond

Example 3.1.5. If \mathbb{k} is any field with an involution $-\circ$, let $L = \mathbb{k}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ be the ring of Laurent polynomials and define the involution on L by

$$\left(\sum_{\alpha} p_{\alpha} x^{\alpha} \right)^{\circ} := \sum_{\alpha} p_{\alpha}^{\circ} x^{-\alpha},$$

where $p_{\alpha} \in \mathbb{k}$, $\alpha \in \mathbb{Z}^n$. Indeed, this turns L into an involutive algebra as the involution is multiplicative since

$$(pq)^{\circ} = \left(\sum_{\alpha, \beta} p_{\alpha} q_{\beta} x^{\alpha + \beta} \right)^{\circ} = \sum_{\alpha, \beta} p_{\alpha}^{\circ} q_{\beta}^{\circ} x^{-(\alpha + \beta)} = p^{\circ} q^{\circ}$$

and satisfies

$$(cp)^{\circ} = \sum_{\alpha} (cp_{\alpha})^{\circ} x^{-\alpha} = c^{\circ} p^{\circ}$$

for $c \in \mathbb{k}$ and Laurent polynomials $p = \sum_{\alpha} p_{\alpha} x^{\alpha}$, $q = \sum_{\beta} q_{\beta} x^{\beta} \in L$. The other requirements clearly hold as well.

For the filtration on R , in this situation we usually pick the one that is induced by max-degree, namely

$$R_{\leq d} = \left\{ \sum_{\alpha} p_{\alpha} x^{\alpha} \in R \mid 0 \leq \alpha_1, \dots, \alpha_n \leq d \right\},$$

since L is the coordinate ring of the algebraic torus.

Of particular interest is the case $\mathbb{k} = \mathbb{C}$ of complex numbers with complex conjugation as involution. In this case, an observation that can be significant in some applications is the following: If we restrict a Laurent polynomial $p \in L$ to the complex torus \mathbb{T}^n , then the involution p° is the complex conjugate of p as a function on \mathbb{T}^n , so we have

$$p^{\circ}(\xi) = \overline{p(\xi)}$$

for all $\xi \in \mathbb{T}^n$, since $\xi^{-\alpha} = \overline{\xi^{\alpha}}$ for all $\alpha \in \mathbb{Z}^n$. In particular, the Laurent polynomial p is a real function on \mathbb{T}^n if and only if $p^{\circ} = p$, i. e. $p_{\alpha} = \overline{p_{-\alpha}}$ for all α . Furthermore, note that, if $\mathfrak{a} \subseteq L$ is a vanishing ideal of a set contained in \mathbb{T}^n , then it follows that $\mathfrak{a}^{\circ} = \mathfrak{a}$. \diamond

3.1 SESQUILINEARITY

Definition 3.1.6. Let $\sigma: L \rightarrow \mathbb{k}$ be a \mathbb{k} -linear map. Then we define the \mathbb{k} -sesquilinear form

$$\langle -, - \rangle_\sigma: L \times L \longrightarrow \mathbb{k}, \quad (q, p) \longmapsto \sigma(q^\circ p),$$

which is $^\circ$ -semilinear in the first and linear in the second argument. Defining sesquilinear forms to be semilinear in the first rather than in the second argument is an arbitrary choice. We choose this convention as it simplifies some notation later on.

By restriction, we can also view this as a sesquilinear form on R as well as on the finite-dimensional vector spaces $R_{\leq d}$, $d \in \mathbb{N}$. Note that this is a symmetric bilinear form if the involution is trivial.

A form $\langle -, - \rangle$ on a \mathbb{k} -vector space U is *Hermitian* if $\langle q, p \rangle = \langle p, q \rangle^\circ$ for all $p, q \in U$. If the involution is trivial, as in Example 3.1.4, then this always holds for $\langle -, - \rangle_\sigma$, as the form is symmetric in that case.

When \mathbb{k} is (a subfield of) the complex numbers \mathbb{C} , then a Hermitian form $\langle -, - \rangle_\sigma$ on U is *positive-semidefinite* if, additionally, $\langle p, p \rangle_\sigma \geq 0$ for all $p \in U$. Note that this never holds if $\mathbb{k} \not\subseteq \mathbb{R}$ and the involution is linear, rather than $^\circ$ -semilinear, unless the form is trivial, since $\langle cp, cp \rangle = c^2 \langle p, p \rangle \geq 0$ cannot hold for all $c \in \mathbb{k}$ and $p \in U$. Thus, one may primarily think of complex conjugation as involution as described in Example 3.1.5 when considering positive-semidefiniteness. In this case, it holds in particular that $\langle cp, cp \rangle = |c|^2 \langle p, p \rangle \geq 0$ for $c \in \mathbb{k} \subseteq \mathbb{C}$ if the form is positive-semidefinite.

Remark 3.1.7. Assume that a family of monomials $\{x^\alpha\}_{\alpha \in J} \subseteq R_{\leq d}$ for a suitable index set $J \subseteq \mathbb{N}^n$ forms a basis of the finite-dimensional vector space $R_{\leq d}$ and that the involution $-\circ$ is trivial. Then the Gramian matrix of $\langle -, - \rangle_\sigma$ with respect to this basis is of the form

$$\left(\langle x^\alpha, x^\beta \rangle_\sigma \right)_{\alpha, \beta \in J} = \left(\sigma((x^\alpha)^\circ x^\beta) \right)_{\alpha, \beta \in J} = \left(\sigma(x^{\alpha+\beta}) \right)_{\alpha, \beta \in J},$$

which is a (generalized) Hankel matrix.

Likewise, if $\{x^\alpha\}_{\alpha \in J} \subseteq R_{\leq d}$ is a basis of $R_{\leq d}$, but L is the ring of Laurent polynomials with involution $-\circ: L \rightarrow L$ defined as in Example 3.1.5, then the Gramian matrix with respect to this basis is a (generalized) Toeplitz matrix of the form

$$\left(\langle x^\alpha, x^\beta \rangle_\sigma \right)_{\alpha, \beta \in J} = \left(\sigma((x^\alpha)^\circ x^\beta) \right)_{\alpha, \beta \in J} = \left(\sigma(x^{-\alpha+\beta}) \right)_{\alpha, \beta \in J}.$$

One may also interpret this matrix in terms of the *bilinear* map

$$R_{\leq d}^\circ \times R_{\leq d} \longrightarrow \mathbb{k}, \quad (q, p) \longmapsto \sigma(qp),$$

by using $\{x^{-\alpha}\}_{\alpha \in J}$ as basis of $R_{\leq d}^\circ$. ◇

Definition 3.1.8. Let V, W be vector spaces over \mathbb{k} and let $\varphi: V \times W \rightarrow \mathbb{k}$ be a sesquilinear map. Then φ is called *right-non-degenerate* if it holds that

$$\{p \in W \mid \varphi(q, p) = 0 \text{ for all } q \in V\} = 0.$$

Similarly, it is *left-non-degenerate* if $\{q \in V \mid \varphi(q, p) = 0 \text{ for all } p \in W\} = 0$ and it is *non-degenerate* if it is both left- and right-non-degenerate.

For future reference, we mention the following fact.

Lemma 3.1.9. *Let V, W be two \mathbb{k} -vector spaces of the same finite dimension and with bases B, B' , respectively. Let $\varphi: V \times W \rightarrow \mathbb{k}$ be a sesquilinear map. Then the following are equivalent:*

- (1) *The matrix $(\varphi(v, w))_{v \in B, w \in B'}$ is invertible.*
- (2) *φ is left-non-degenerate.*
- (3) *φ is right-non-degenerate.*

Proof. As V and W are finite-dimensional, the map φ is right-non-degenerate (left-non-degenerate) if and only if the matrix $(\varphi(v, w))_{v \in B, w \in B'}$ has full column (row) rank. As V and W have the same dimension, this is equivalent to the non-singularity of the matrix. \square

Lemma 3.1.10. *Let $U \subseteq L$ be a \mathbb{k} -vector subspace and let $\sigma: L \rightarrow \mathbb{k}$ be a \mathbb{k} -linear map. Then the sesquilinear form $\langle -, - \rangle_\sigma: U \times U \rightarrow \mathbb{k}$ defined by $\langle q, p \rangle_\sigma = \sigma(q^\circ p)$ is Hermitian if and only if $\sigma(g^\circ) = \sigma(g)^\circ$ for all $g \in U^\circ \cdot U$.*

If $U = L$, note in particular that $U^\circ \cdot U = L$.

Proof. If the form is Hermitian, write $g \in U^\circ \cdot U$ as $g = q^\circ p$ for suitable $q, p \in U$. Then we have $\sigma(g^\circ) = \sigma(p^\circ q) = \langle p, q \rangle_\sigma = \langle q, p \rangle_\sigma^\circ = \sigma(q^\circ p)^\circ = \sigma(g)^\circ$, as the form is Hermitian. Conversely, if $q, p \in U$, then $g = q^\circ p \in U^\circ \cdot U$, so it follows in a similar manner that the form is Hermitian. \square

Lemma 3.1.11. *Assume that $\sigma: L \rightarrow \mathbb{k}$ is a \mathbb{k} -linear map, $\mathfrak{a} \subseteq L$ is an ideal such that $\mathfrak{a}, \mathfrak{a}^\circ \subseteq \ker \sigma$. Then the sesquilinear form $\langle -, - \rangle_\sigma$ on L induces a sesquilinear form*

$$R_{\leq d}/(\mathfrak{a} \cap R_{\leq d}) \times R_{\leq d}/(\mathfrak{a} \cap R_{\leq d}) \longrightarrow \mathbb{k}, \quad (\bar{q}, \bar{p}) \longmapsto \langle q, p \rangle_\sigma = \sigma(q^\circ p),$$

for any $d \in \mathbb{N}$.

Here, \bar{q}, \bar{p} denotes the residue class of polynomials $q, p \in R_{\leq d}$ modulo $\mathfrak{a} \cap R_{\leq d}$. We denote the induced sesquilinear form on $R_{\leq d}/(\mathfrak{a} \cap R_{\leq d})$ by $\langle -, - \rangle_\sigma$ again. Also note that it follows from Lemma 3.1.10 that the requirements $\mathfrak{a} \subseteq \ker \sigma$ and $\mathfrak{a}^\circ \subseteq \ker \sigma$ are equivalent when the sesquilinear form $\langle -, - \rangle_\sigma$ on L is Hermitian.

Proof. Let $p, q \in R_{\leq d}$. If $p \in \mathfrak{a} \cap R_{\leq d}$, then $q^\circ p$ is contained in $\mathfrak{a} \subseteq \ker \sigma$, so $\sigma(q^\circ p) = 0$. Likewise, if $q \in \mathfrak{a} \cap R_{\leq d}$, then $q^\circ p \in \mathfrak{a}^\circ \subseteq \ker \sigma$, so the sesquilinear form on $R_{\leq d}/(\mathfrak{a} \cap R_{\leq d})$ is well-defined. \square

Definition 3.1.12. If $\sigma: L \rightarrow \mathbb{k}$ is \mathbb{k} -linear and $\mathfrak{a} \subseteq L$ is an ideal such that $\mathfrak{a} \subseteq \ker \sigma$, then the sesquilinear form $\langle -, - \rangle_\sigma$ on L does *not* induce a sesquilinear form on the quotient spaces $R_{\leq d}/(\mathfrak{a} \cap R_{\leq d})$, $d \leq \mathbb{N}$. (Observe that this would need $\mathfrak{a}^\circ \subseteq \ker \sigma$ or require the

3.1 SESQUILINEARITY

form to be Hermitian, as in Lemma 3.1.11.) However, for every vector subspace $U \subseteq L$, thus in particular for $U = R_{\leq d}$, we can still define the sesquilinear map

$$\Phi_{\mathfrak{a},\sigma}^U: U/(\mathfrak{a}^\circ \cap U) \times U/(\mathfrak{a} \cap U) \longrightarrow \mathbb{k}, \quad (\bar{q}, \bar{p}) \longmapsto \langle q, p \rangle_\sigma = \sigma(q^\circ p),$$

which is well-defined in the semilinear argument since $\sigma(q^\circ p) = 0$ for any $p \in U$ and $q \in \mathfrak{a}^\circ \cap U$, as $q^\circ p \in \mathfrak{a}$. It is important not to confuse this map $\Phi_{\mathfrak{a},\sigma}^U$ with a potential sesquilinear form on $U/(\mathfrak{a} \cap U)$ as given in Lemma 3.1.11. If the involution is trivial, then these two sesquilinear maps are the same, but in general they are distinct and can have quite different properties. Although there exist non-trivial involutions such that the quotient spaces $U/(\mathfrak{a}^\circ \cap U)$ and $U/(\mathfrak{a} \cap U)$ have the same dimension, as in Example 3.1.13, in general this is not the case, so they can be of different dimensions; see for instance Example 3.1.15. Nevertheless, we can relate these two sesquilinear maps by the notion of right-non-degenerateness, as is expressed in Lemma 3.1.16.

In the remainder, we refer to $\Phi_{\mathfrak{a},\sigma}^U$ as the *induced sesquilinear map* associated to σ or as the map induced by the sesquilinear form $\langle -, - \rangle_\sigma$ on L . If σ is clear from the context, we may also denote the map by $\Phi_{\mathfrak{a}}^U$ or simply by Φ .

Example 3.1.13. Let $R = \mathbb{k}[x_1, \dots, x_n]$ and $L = \mathbb{k}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ be the polynomial and Laurent polynomial rings and define an involution on L as in Example 3.1.5. Moreover, let $U := \bigoplus_{\alpha \in \mathbb{N}^n, \alpha_i \leq \beta_i} \mathbb{k}x^\alpha \subseteq R$ for some $\beta = (\beta_1, \dots, \beta_n) \in \mathbb{N}^n$. Then $U^\circ = x^{-\beta}U$, so U and U° have the same dimension. The map $L \rightarrow L, p \mapsto x^{-\beta}p$, is an automorphism of L under which \mathfrak{a} is invariant, as $x^{-\beta}$ is a unit in L . Therefore, a polynomial p is contained in $\mathfrak{a} \cap U$ if and only if the Laurent polynomial $x^{-\beta}p$ is contained in $\mathfrak{a} \cap U^\circ$, so we conclude that $\mathfrak{a} \cap U$ and $\mathfrak{a}^\circ \cap U$ have the same dimension. \diamond

Remark 3.1.14. By a similar argument as in Example 3.1.13, one can observe that an ideal \mathfrak{a} in L is generated by Laurent polynomials $p \in \mathfrak{a}$ bounded in max-degree $d \in \mathbb{N}$ that are of the form $p = \sum_{\alpha \in \mathbb{Z}^n, |\alpha|_\infty \leq d} p_\alpha x^\alpha$, $p_\alpha \in \mathbb{k}$, if and only if it is generated by polynomials $\mathfrak{a} \cap R_{\leq 2d}$, where $R_{\leq 2d}$ denotes a filtered component with respect to the filtration on R that is induced by max-degree. \diamond

Example 3.1.15. Let $R = \mathbb{k}[x]$ and $L = \mathbb{k}[x, y]$. Define the involution on L by $x^\circ = y$ and $y^\circ = x$, extending the trivial involution on \mathbb{k} . Further, let $\mathfrak{a} = \langle x \rangle \subseteq L$ and $U = R_{\leq d}$ for some $d \geq 1$. Then the ideal $\mathfrak{a}^\circ = \langle y \rangle$ intersects trivially with U , so we have $\dim(\mathfrak{a}^\circ \cap U) = 0$, while $\dim(\mathfrak{a} \cap U) > 0$. Consequently, a sesquilinear map $U/(\mathfrak{a}^\circ \cap U) \times U/(\mathfrak{a} \cap U) \rightarrow \mathbb{k}$ cannot be non-degenerate. \diamond

Lemma 3.1.16. *Let $\sigma: L \rightarrow \mathbb{k}$ be a \mathbb{k} -linear map, let $\mathfrak{a} \subseteq L$ be an ideal satisfying $\mathfrak{a}, \mathfrak{a}^\circ \subseteq \ker \sigma$ and let $U \subseteq L$ be a vector subspace. Denote by ψ, Φ the sesquilinear maps*

$$\psi: U/(\mathfrak{a} \cap U) \times U/(\mathfrak{a} \cap U) \rightarrow \mathbb{k}, \quad \Phi: U/(\mathfrak{a}^\circ \cap U) \times U/(\mathfrak{a} \cap U) \rightarrow \mathbb{k}$$

that are induced by the sesquilinear form $\langle -, - \rangle_\sigma$ on L . Then ψ is right-non-degenerate if and only if Φ is right-non-degenerate.

Note that, if U is finite-dimensional and ψ is right-non-degenerate, then it is already non-degenerate by Lemma 3.1.9.

Proof. Denote by π, π' the natural quotient maps from U to $U/(\mathfrak{a} \cap U)$ and $U/(\mathfrak{a}^\circ \cap U)$, respectively. As both Φ and ψ are induced by $\langle -, - \rangle_\sigma$, we have

$$\Phi(\pi'(q), \pi(p)) = \langle q, p \rangle_\sigma = \psi(\pi(q), \pi(p))$$

for all $p, q \in U$. Thus, it follows from the definition that Φ is right-non-degenerate if and only if ψ is right-non-degenerate. \square

We frequently make use of the following simple and useful correspondence. It allows us to switch context from degenerate sesquilinear maps in the ambient ring to right-non-degenerate sesquilinear maps on the quotient ring modulo an ideal. Primarily, this is needed so that involved matrices form a matrix pencil that is regular, which is a requirement for the study of eigenvalues. This is explained in detail in Section 4.2.

Lemma 3.1.17. *Let $\sigma: L \rightarrow \mathbb{k}$ be a \mathbb{k} -linear map and let $\mathfrak{a} \subseteq L$ be an ideal such that $\mathfrak{a} \subseteq \ker \sigma$. Let $U \subseteq R$ be a \mathbb{k} -vector subspace and assume that the induced sesquilinear map $\Phi_{\mathfrak{a}, \sigma}^U$ is right-non-degenerate. Let $p \in U$ be a polynomial. Then the following are equivalent:*

- (1) $p \in \mathfrak{a}$;
- (2) $\sigma(q^\circ p) = 0$ for all $q \in U$;
- (3) $H(p) = 0$, where H is the \mathbb{k} -linear map

$$H: U \longrightarrow \text{Hom}_{\mathbb{k}}^{\text{semi}}(U, \mathbb{k}), \quad p \longmapsto (q \mapsto \sigma(q^\circ p)).$$

Here, $\text{Hom}_{\mathbb{k}}^{\text{semi}}(U, \mathbb{k})$ denotes the vector space of $^\circ$ -semilinear maps from U to \mathbb{k} .

Proof. It is clear that (2) and (3) are equivalent. Moreover, if $p \in \mathfrak{a}$, we also have $q^\circ p \in \mathfrak{a} \subseteq \ker \sigma$ for all $q \in L$, thus in particular for all $q \in U \subseteq R \subseteq L$, so (1) implies (2). Conversely, assume that $\sigma(q^\circ p) = 0$ for all $q \in U$. As $\mathfrak{a} \subseteq \ker \sigma$, we have

$$\Phi_{\mathfrak{a}, \sigma}^U(q + \mathfrak{a}^\circ \cap U, p + \mathfrak{a} \cap U) = \sigma((q^\circ + \mathfrak{a} \cap U^\circ)(p + \mathfrak{a} \cap U)) = \sigma(q^\circ p) = 0$$

for all $q \in U$. Since the sesquilinear map $\Phi_{\mathfrak{a}, \sigma}^U$ is right-non-degenerate, it follows that $p \equiv 0 \pmod{\mathfrak{a} \cap U}$ and thus $p \in \mathfrak{a}$. \square

3.2 Factorization properties

The Vandermonde factorization of Lemma 1.3.3 is an essential aspect of Prony's method. Here, we analyze how to transfer it from measures on zero-dimensional to measures on positive-dimensional algebraic varieties. The statements here are also motivated by the study of finite-rank Hankel operators as in e.g. [Mou18]. In the positive-dimensional

setting, such operators are not of finite rank anymore, but some properties are still valid.

Let \mathbb{k}, R, L be as in Section 3.1, so \mathbb{k} is a field of characteristic 0, $R = \mathbb{k}[x_1, \dots, x_n]$ is the polynomial ring in n variables endowed with a filtration $\{R_{\leq d}\}_{d \in \mathbb{N}}$ and L is an involutive commutative \mathbb{k} -algebra such that $R \subseteq L$.

We wish to examine more closely the following situation. Let $\mathfrak{a} \subseteq L$ be an ideal and let $\sigma: L \rightarrow \mathbb{k}$ be a \mathbb{k} -linear map with the property that $\mathfrak{a} \subseteq \ker \sigma$. This means that the map σ factors via the quotient homomorphism

$$\pi_{\mathfrak{a}}: L \longrightarrow L/\mathfrak{a}, \quad p \longmapsto \bar{p} := p + \mathfrak{a},$$

which we denote by $\pi_{\mathfrak{a}}$, and a \mathbb{k} -linear map $\bar{\sigma}: L/\mathfrak{a} \rightarrow \mathbb{k}$, denoted by $\bar{\sigma}$.

Example 3.2.1. Assume that L is the polynomial ring R and $\xi \in \mathbb{k}^n$ (or that L is the Laurent polynomial ring in n variables and $\xi \in (\mathbb{k}^*)^n$). Then, for the maximal ideal $\mathfrak{m}_{\xi} = \langle x - \xi \rangle \subseteq L$, this gives the evaluation homomorphism at the point ξ ,

$$\pi_{\mathfrak{m}_{\xi}}: L \longrightarrow L/\mathfrak{m}_{\xi} \cong \mathbb{k}, \quad x^{\alpha} \longmapsto \bar{x}^{\alpha} = \xi^{\alpha},$$

for $\alpha \in \mathbb{N}^n$ (or $\alpha \in \mathbb{Z}^n$), so $\pi_{\mathfrak{m}_{\xi}}(p) = p(\xi)$ for $p \in L$. Note further that, for any \mathbb{k} -linear map $\sigma: L \rightarrow \mathbb{k}$ with $\mathfrak{m}_{\xi} \subseteq \ker \sigma$, the linear map $\bar{\sigma}: L/\mathfrak{m}_{\xi} \cong \mathbb{k} \rightarrow \mathbb{k}$ is determined by a single scalar $\lambda \in \mathbb{k}$, with respect to a suitable basis. Thus, $\sigma = \lambda \pi_{\mathfrak{m}_{\xi}} = \lambda \text{ev}_{\xi} \in \text{Hom}_{\mathbb{k}}(L, \mathbb{k})$, which we can interpret as an exponential sum of rank 1 if $\lambda \neq 0$ (cf. Section 1.3.1).

More generally, consider the zero-dimensional ideal $\mathfrak{a} = \bigcap_{j=1}^r \mathfrak{m}_{\xi_j}$, for distinct points ξ_1, \dots, ξ_r . Then it follows from Lemma 1.2.2 that

$$L/\mathfrak{a} \cong \bigoplus_{j=1}^r L/\mathfrak{m}_{\xi_j} \cong \mathbb{k}^r$$

and $\pi_{\mathfrak{a}}(p) = (p(\xi_1), \dots, p(\xi_r))$ for $p \in L$. As a \mathbb{k} -linear map with respect to the monomial basis of L , we can view this map as being described by an infinite Vandermonde matrix associated to the points ξ_1, \dots, ξ_r . If $\sigma: L \rightarrow \mathbb{k}$ is a \mathbb{k} -linear map with $\mathfrak{a} \subseteq \ker \sigma$, then it is of the form $\sigma = \sum_{j=1}^r \lambda_j \text{ev}_{\xi_j}$ with suitable parameters $\lambda_1, \dots, \lambda_r \in \mathbb{k}$, which corresponds to an exponential sum of rank r if $\lambda_1, \dots, \lambda_r \neq 0$. \diamond

The following example shows that the ideal \mathfrak{a} does not need to be radical, which does not classically correspond to an algebraic variety.

Example 3.2.2. Let $n = 1$ and $L = R = \mathbb{k}[x]$ be the univariate polynomial ring. Let $\xi \in \mathbb{k}$ and define the primary ideal $\mathfrak{a} = \mathfrak{m}_{\xi}^2 = \langle x - \xi \rangle^2$. Then

$$\begin{aligned} \pi_{\mathfrak{a}}: L &\longrightarrow L/\mathfrak{a} \cong \mathbb{k}1 \oplus \mathbb{k}(\bar{x} - \xi), \\ x^{\alpha} &\longmapsto \bar{x}^{\alpha} = \xi^{\alpha} + \alpha \xi^{\alpha-1}(\bar{x} - \xi), \end{aligned}$$

for $\alpha \in \mathbb{N}$. Let $\sigma: L \rightarrow \mathbb{k}$ be a \mathbb{k} -linear map with $\mathfrak{a} \subseteq \ker \sigma$. Then $\bar{\sigma}$ is determined by the images of the basis elements $1, \bar{x} - \xi$. Assume that $\bar{\sigma}(1) = \lambda$, $\bar{\sigma}(\bar{x} - \xi) = \lambda'$ for some $\lambda, \lambda' \in \mathbb{k}$. Then

$$\sigma(x^\alpha) = \lambda \xi^\alpha + \lambda' \alpha \xi^{\alpha-1} = \text{ev}_\xi((\lambda + \lambda' \partial)(x^\alpha)),$$

for $\alpha \in \mathbb{N}$, so σ is a moment functional of a first-order local mixture of a Dirac distribution, as studied in Chapter 2. More generally, this is an example of a polynomial-exponential series; see for instance [Mou18]. \diamond

As R is endowed with a filtration $\{R_{\leq d}\}_{d \in \mathbb{N}}$ for which each component $R_{\leq d}$ is finite-dimensional and since $R \subseteq L$, we can restrict the map $\pi_{\mathfrak{a}}: L \rightarrow L/\mathfrak{a}$ to a map on finite-dimensional vector subspaces $R_{\leq d} \rightarrow R_{\leq d}/(\mathfrak{a} \cap R_{\leq d})$, which we denote by $\pi_{\mathfrak{a}, \leq d}$, as explained in Example 3.1.3.

Remark 3.2.3. The isomorphism $L/\mathfrak{a} \cong \bigoplus_{j=1}^r L/\mathfrak{m}_{\xi_j}$ in Example 3.2.1 only holds because the ideals $\mathfrak{m}_{\xi_1}, \dots, \mathfrak{m}_{\xi_r}$ are pairwise comaximal. For general ideals $\mathfrak{a}_1, \dots, \mathfrak{a}_r$ with $\mathfrak{a} = \bigcap_{j=1}^r \mathfrak{a}_j$, this does not hold.

Even if the ideals $\mathfrak{a}_1, \dots, \mathfrak{a}_r$ are pairwise comaximal, it is not in general possible to find an isomorphism between the truncated spaces $R_{\leq d}/(\mathfrak{a} \cap R_{\leq d})$ and $\bigoplus_{j=1}^r R_{\leq d}/(\mathfrak{a}_j \cap R_{\leq d})$, for any $d \in \mathbb{N}$. This suggests that it is often more useful to consider $R_{\leq d}/(\mathfrak{a} \cap R_{\leq d})$ rather than $\bigoplus_{j=1}^r R_{\leq d}/(\mathfrak{a}_j \cap R_{\leq d})$, even though, for any d , the latter space can be interpreted as the codomain of the map $\text{ev}_{\leq d}$ that appears in Theorem 1.3.1. \diamond

An important ingredient of Prony's method is that we can extract information about the vanishing ideal from the kernel of the moment matrix, if the moment matrix is sufficiently large; see Theorems 1.3.1 and 1.3.6 and Lemma 1.3.3. In the following, we examine what is required to transfer this property to the setting of ideals which are possibly not of dimension zero, but are of higher dimension. This is answered by the following theorem as well as Corollary 3.2.5 below.

Theorem 3.2.4. *Let $\mathfrak{a} \subseteq L$ be an ideal and let $\sigma: L \rightarrow \mathbb{k}$ be a \mathbb{k} -linear map with $\mathfrak{a} \subseteq \ker \sigma$. Then the \mathbb{k} -linear map*

$$H: R \longrightarrow \text{Hom}_{\mathbb{k}}^{\text{semi}}(R, \mathbb{k}), \quad p \longmapsto (q \mapsto \langle q, p \rangle_\sigma),$$

factors as

$$\begin{aligned} R &\xrightarrow{\pi_{\mathfrak{a}}} R/(\mathfrak{a} \cap R) \longrightarrow \text{Hom}_{\mathbb{k}}^{\text{semi}}(R/(\mathfrak{a}^\circ \cap R), \mathbb{k}) \xrightarrow{\pi_{\mathfrak{a}^\circ}^\top} \text{Hom}_{\mathbb{k}}^{\text{semi}}(R, \mathbb{k}), \\ p + \mathfrak{a} \cap R &\longmapsto (q + \mathfrak{a}^\circ \cap R \mapsto \langle q, p \rangle_\sigma), \end{aligned}$$

where $\pi_{\mathfrak{a}^\circ}^\top(\varphi) = \varphi \circ \pi_{\mathfrak{a}^\circ}$ for $\varphi \in \text{Hom}_{\mathbb{k}}^{\text{semi}}(R/(\mathfrak{a}^\circ \cap R), \mathbb{k})$.

Moreover, the truncated map between finite-dimensional vector subspaces given by

$$H_{d+\delta, d}: R_{\leq d} \longrightarrow \text{Hom}_{\mathbb{k}}^{\text{semi}}(R_{\leq d+\delta}, \mathbb{k}), \quad p \longmapsto (q \mapsto \langle q, p \rangle_\sigma),$$

for $d, \delta \in \mathbb{N}$, factors as

$$\begin{array}{ccc}
 R_{\leq d} & \xrightarrow{H_{d+\delta,d}} & \text{Hom}_{\mathbb{k}}^{\text{semi}}(R_{\leq d+\delta}, \mathbb{k}) \\
 \downarrow \pi_{\mathfrak{a}, \leq d} & & \uparrow \pi_{\mathfrak{a}^\circ, \leq d+\delta}^\top \\
 R_{\leq d}/(\mathfrak{a} \cap R_{\leq d}) & \xrightarrow{\overline{H_{d+\delta,d}}} & \text{Hom}_{\mathbb{k}}^{\text{semi}}(R_{\leq d+\delta}/(\mathfrak{a}^\circ \cap R_{\leq d+\delta}), \mathbb{k}), \\
 p + \mathfrak{a} \cap R_{\leq d} & \longmapsto & (q + \mathfrak{a}^\circ \cap R_{\leq d+\delta} \mapsto \langle q, p \rangle_\sigma).
 \end{array} \tag{3.1}$$

Proof. Due to the inclusion $\mathfrak{a} \subseteq \ker \sigma$, we have that

$$\sigma((q + \mathfrak{a}^\circ \cap R)^\circ(p + \mathfrak{a} \cap R)) = \sigma((q^\circ + \mathfrak{a} \cap R^\circ)(p + \mathfrak{a} \cap R)) = \sigma(q^\circ p) = \langle q, p \rangle_\sigma,$$

for all $q, p \in R$, which shows the first factorization property. The other one follows analogously. \square

The truncated map $H_{d+\delta,d}$ is of importance for us, since we are interested in recovery from finitely many moments. By Theorem 3.2.4, it holds that

$$\mathfrak{a} \cap R_{\leq d} \subseteq \ker H_{d+\delta,d}$$

and we ask when this is an equality. This leads to the following corollary.

Corollary 3.2.5. *If the map $\overline{H_{d+\delta,d}}: R_{\leq d}/(\mathfrak{a} \cap R_{\leq d}) \rightarrow \text{Hom}_{\mathbb{k}}^{\text{semi}}(R_{\leq d+\delta}/(\mathfrak{a}^\circ \cap R_{\leq d+\delta}), \mathbb{k})$ is injective, then*

$$\ker(H_{d+\delta,d}) = \ker(\pi_{\mathfrak{a}, \leq d}) = \mathfrak{a} \cap R_{\leq d}.$$

Proof. Due to the factorization (3.1) and since the map $\pi_{\mathfrak{a}^\circ, \leq d+\delta}^\top$ is injective, the equality holds if and only if the map $\overline{H_{d+\delta,d}}$ is injective. \square

As the vector space dimension of the codomain of $\overline{H_{d+\delta,d}}$ is finite and at least as large as the dimension of the domain, saying that $\overline{H_{d+\delta,d}}$ is injective is the same as saying that the map $\overline{H_{d+\delta,d}}$ has full rank. As such, this can be regarded as a variant of the statement about the Vandermonde factorization in Lemma 1.3.3.

Remark 3.2.6. In this formalism, $\pi_{\mathfrak{a}, \leq d}$ is always surjective, which is an important difference from the Vandermonde factorization considered in Lemma 1.3.3, as the Vandermonde matrix considered there can be non-surjective for small d . This is explained further in Example 3.2.7 below. There, for an ideal of the form $\mathfrak{a} = \bigcap_{j=1}^r \mathfrak{m}_{\xi_j}$, the dimension of $\text{im}(\pi_{\mathfrak{a}, \leq d}) = R_{\leq d}/(\mathfrak{a} \cap R_{\leq d})$ as vector space is at most r , but can be smaller. Equality holds if and only if the corresponding Vandermonde matrix has rank r , which only holds if d is sufficiently large. \diamond

Note that $\overline{H_{d+\delta,d}}$ has full rank if and only if it induces a right-non-degenerate sesquilinear map of the form

$$R_{\leq d+\delta}/(\mathfrak{a}^\circ \cap R_{\leq d+\delta}) \times R_{\leq d}/(\mathfrak{a} \cap R_{\leq d}) \longrightarrow \mathbb{k}, \quad (q, p) \longmapsto \overline{H_{d+\delta,d}}(p)(q).$$

Whether this is the case depends on the functional σ . We will see in Theorem 3.4.11 that this holds in particular when σ is a moment functional of a measure and \mathfrak{a} is the vanishing ideal of its support.

Example 3.2.7. Let us revisit Example 3.2.1, so let $\mathfrak{a} := \bigcap_{j=1}^r \mathfrak{m}_{\xi_j} \subseteq L$ for distinct points $\xi_1, \dots, \xi_r \in \mathbb{k}^n$, where now we assume that $L = R = \mathbb{k}[x_1, \dots, x_n]$ is endowed with the trivial involution and the filtration induced by total degree.

If d is sufficiently large, $\pi_{\mathfrak{a}, \leq d}$ has rank r and we have $R_{\leq d}/(\mathfrak{a} \cap R_{\leq d}) \cong \bigoplus_{j=1}^r R/\mathfrak{m}_{\xi_j} \cong \mathbb{k}^r$. Hence, we also have $R_{\leq d+\delta}/(\mathfrak{a} \cap R_{\leq d+\delta}) \cong \mathbb{k}^r$ for all $\delta \in \mathbb{N}$. If $\sigma: R \rightarrow \mathbb{k}$ is a \mathbb{k} -linear map with $\mathfrak{a} \subseteq \ker \sigma$, then, by Example 3.2.1, it is of the form $\sigma = \sum_{j=1}^r \lambda_j \text{ev}_{\xi_j}$ for some $\lambda_1, \dots, \lambda_r \in \mathbb{k}$. Thus, the map $\overline{H_{d+\delta, d}}$ corresponds to the diagonal matrix $\text{diag}(\lambda_1, \dots, \lambda_r)$ with respect to the natural bases. Clearly, it is injective if and only if $\lambda_1, \dots, \lambda_r \neq 0$, which illustrates the connection of Corollary 3.2.5 to Lemma 1.3.3. \diamond

Although for zero-dimensional ideals as in the preceding example it is enough to consider the case $\delta = 0$ to infer that $\ker H_{d+\delta, d} = \mathfrak{a} \cap R_{\leq d}$ if d is sufficiently large, this does not hold in general (cf. Example 3.4.13). We will see a non-trivial example in Example 3.4.9, which involves an ideal of positive dimension. In connection to that, Theorem 3.4.3 will show that it can be useful to consider δ larger than 0.

3.3 Considerations for the complex torus

Here we briefly examine the complex torus from an algebraic point of view and list some implications for the moment problems we are interested in. In particular, this helps understand the relevance of Hankel and Toeplitz matrices. See also [Sch17, Chapters 15.1, 15.2] for a related, more extensive treatment.

The complex torus $\mathbb{T}^1 = \{z \in \mathbb{C} \mid |z| = 1\} \subseteq \mathbb{C}$ is not a complex algebraic variety. It is a Zariski-dense set in \mathbb{C} , so its Zariski closure is the whole space \mathbb{C} . However, it can be viewed as the real algebraic variety with coordinate ring $S := \mathbb{R}[x, y]/\langle x^2 + y^2 - 1 \rangle$. Over the complex numbers, we have the isomorphism

$$\begin{aligned} \mathbb{C}[z, z^{-1}] = \mathbb{C}[z, z']/\langle zz' - 1 \rangle &\xrightarrow{\sim} \mathbb{C}[x, y]/\langle x^2 + y^2 - 1 \rangle = S \otimes_{\mathbb{R}} \mathbb{C}, \\ (z, z') &\longmapsto (x + iy, x - iy). \end{aligned} \tag{3.2}$$

Over the reals, this is not possible. Set $u := (\frac{1}{2}(z + z^{-1}), \frac{1}{2i}(z - z^{-1}))$ and identify it with the residue class of (x, y) . We have the following simple fact, which says that this transformation is degree compatible. In other words, we obtain a morphism of filtered algebras.

Lemma 3.3.1. *For $d \in \mathbb{N}$, the subspaces $\langle z^{-d}, \dots, z^d \rangle$ and $\langle u^\alpha \mid \alpha \in \mathbb{N}^2, |\alpha| \leq d \rangle$ of $\mathbb{C}[z^{\pm 1}]$ are equal.*

Proof. Observe that $u^\alpha = \varphi^{-1}((x, y)^\alpha)$ involves only the monomials z^{-d}, \dots, z^d if $|\alpha| \leq d$. \square

This means that, for a moment functional $\sigma: \mathbb{C}[z^{\pm 1}] \rightarrow \mathbb{C}$, the moments $\sigma(z^{-d}), \dots, \sigma(z^d)$ contain the same information as the moments $\sigma(u^\alpha) = \sigma(\varphi^{-1}(x^{\alpha_1} y^{\alpha_2}))$, $\alpha \in \mathbb{N}^2$, $|\alpha| \leq d$, and can be converted using the transformation φ . Analogously, Lemma 3.3.1 also holds for the higher-dimensional torus \mathbb{T}^n by tensorization.

Remark 3.3.2. We define an involution $-^\circ$ on $\mathbb{C}[z^{\pm 1}]$ as in Example 3.1.5 by $z^\circ := z^{-1} = \varphi^{-1}(x - iy)$, extending the complex conjugation on \mathbb{C} . Then $-^\circ$ acts trivially on the coordinates u , so $u_1^\circ = \varphi^{-1}(x) = u_1$ and $u_2^\circ = \varphi^{-1}(y) = u_2$.

If $\sigma: \mathbb{C}[z^{\pm 1}] \rightarrow \mathbb{C}$ is a \mathbb{C} -linear map such that the associated sesquilinear form $\langle -, - \rangle_\sigma$ is positive-semidefinite, then, for $d \in \mathbb{N}$, the Toeplitz matrices $T_d := (\langle z^\alpha, z^\beta \rangle_\sigma)_{0 \leq \alpha, \beta \leq d} = (\sigma(z^{-\alpha+\beta}))_{0 \leq \alpha, \beta \leq d}$ and $T' := (\langle z^\alpha, z^\beta \rangle_\sigma)_{-d \leq \alpha, \beta \leq d} = (\sigma(z^{-\alpha+\beta}))_{-d \leq \alpha, \beta \leq d}$ are positive-semidefinite. They are the Gramian matrices of the form restricted to the subspaces generated by the monomials $1, z, \dots, z^d$ and z^{-d}, \dots, z^d , respectively, as discussed in Remark 3.1.7. Observe that the Toeplitz matrix T' is equal to the Toeplitz matrix $T_{2d} = (\sigma(z^{-\alpha+\beta}))_{0 \leq \alpha, \beta \leq 2d}$. Hence, considering Toeplitz matrices with respect to the monomials $1, \dots, z^d$, for $d \in \mathbb{N}$, leads to a finer filtration than with the monomials z^{-d}, \dots, z^d .

Furthermore, the Hankel matrix $H := (\langle u^\alpha, u^\beta \rangle_\sigma)_{\alpha, \beta \in \mathbb{N}^2, |\alpha|, |\beta| \leq d} = (\sigma(u^{\alpha+\beta}))_{|\alpha|, |\beta| \leq d}$ is positive-semidefinite. Note that the elements u^α , $|\alpha| \leq d$, are not linearly independent, but generate $\langle z^{-d}, \dots, z^d \rangle$ by Lemma 3.3.1, and, for every element $p = \sum_{|\alpha| \leq d} p_\alpha u^\alpha$, $p_\alpha \in \mathbb{C}$, we have

$$(p_\alpha)_{|\alpha| \leq d}^* H (p_\beta)_{|\beta| \leq d} = \sum_{|\alpha|, |\beta| \leq d} \overline{p_\alpha} p_\beta \langle u^\alpha, u^\beta \rangle_\sigma = \langle p, p \rangle_\sigma \geq 0.$$

Since H is Hermitian, but also symmetric, this implies in particular that H must already be a real matrix.

In summary, this establishes a natural relationship between Toeplitz and Hankel matrices associated to a sesquilinear form $\langle -, - \rangle_\sigma$. Similarly, this works for the higher-dimensional torus, as well. For another approach relating Toeplitz and Hankel matrices, see also [KRvdO20] and [vdOhe21, Section 2.1]. \diamond

Example 3.3.3. Let us consider an example on \mathbb{T}^2 where we identify $\mathbb{C}[z_1^{\pm 1}, z_2^{\pm 1}]$ with the complexification of $S \otimes S$. We write $x_1 := x \otimes 1$, $y_1 := y \otimes 1$, $x_2 := 1 \otimes x$, $y_2 := 1 \otimes y$ for the elements in $S \otimes S$ and identify $z_k = x_k + iy_k$, $z_k^{-1} = x_k - iy_k$ for $k = 1, 2$. Define the Laurent polynomial

$$\begin{aligned} f &= z_1^{-1} z_2^{-1} + 2z_1^{-1} + z_1^{-1} z_2 + 2z_2^{-1} - 1 + 2z_2 + z_1 z_2^{-1} + 2z_1 + z_1 z_2 \\ &= (z_1 + 1)(z_1^{-1} + 1)(z_2 + 1)(z_2^{-1} + 1) - 5. \end{aligned} \quad (3.3)$$

Then $f = 4x_1 x_2 + 4x_1 + 4x_2 - 1$. In particular, the polynomial f has max-degree 1 in the variables z_1, z_2 and has bidegree $(1, 1)$ in $(x_1, y_1), (x_2, y_2)$, has real coefficients,

so is contained in $S \otimes S$, and cuts out a real variety on the torus \mathbb{T}^2 . By using the parametrization of the torus $z_k = e^{2\pi i t_k}$, $x_k = \cos(2\pi t_k)$, $y_k = \sin(2\pi t_k)$ for $t_k \in [0, 1)$, $k = 1, 2$, one can visualize this trigonometric curve. It is one of the curves depicted in Figure 4.2 in Section 4.8.3.

Assume that $\sigma: \mathbb{C}[z_1^{\pm 1}, z_2^{\pm 1}] \rightarrow \mathbb{C}$ is the moment functional of the uniform measure (or any other measure with sufficiently large support) supported on this trigonometric curve in \mathbb{T}^2 . Then the Toeplitz moment matrix

$$\left(\sigma(z^{-\gamma+\gamma'}) \right)_{\gamma, \gamma' \in \mathbb{Z}^2, |\gamma|_\infty, |\gamma'|_\infty \leq 1}$$

is positive-semidefinite and its kernel is spanned by the polynomial (3.3). Similarly, the kernel of the Hankel moment matrix

$$\left(\sigma((x_1, y_1, x_2, y_2)^{\alpha+\beta}) \right)_{\alpha, \beta \in \mathbb{N}^2 \times \mathbb{N}^2, |\alpha_1|, |\alpha_2| \leq 1, |\beta_1|, |\beta_2| \leq 1}$$

contains the polynomial $f = 4x_1x_2 + 4x_1 + 4x_2 - 1$. In fact, as the degree is small, f spans the kernel, but for larger Hankel matrices one would also find elements from the ideal $\langle x_1^2 + y_1^2 - 1, x_2^2 + y_2^2 - 1 \rangle$ defining $S \otimes S$ in the kernel. This can be avoided by considering the submatrix corresponding to a linearly independent subset of the monomials.

It is important to note that, even though we might work with complex-valued moments $\sigma(z^\gamma) \in \mathbb{C}$, $\gamma \in \mathbb{Z}^2$, the variety we are dealing with is a real variety defined in terms of the real ring $S \otimes S \cong \mathbb{R}[x_1, y_1, x_2, y_2] / \langle x_1^2 + y_1^2 - 1, x_2^2 + y_2^2 - 1 \rangle$. This ring is not isomorphic to $\mathbb{R}[z_1^{\pm 1}, z_2^{\pm 1}]$. \diamond

3.4 Recovery of the support from moments

In this section, we explore how to recover the underlying algebraic variety that a measure is supported on, by using finitely many of its moments. We consider a non-negative or signed measure μ whose support lives in the affine space \mathbb{R}^n or the complex torus \mathbb{T}^n and wish to find the smallest variety that contains the support. Following the notation of Section 3.1, we consider the following two cases, to which we also refer as *affine* and *trigonometric* cases, respectively:

- (1) $\Omega = \mathbb{R}^n$, $\mathbb{k} = \mathbb{R}$, $L = R = \mathbb{R}[x_1, \dots, x_n]$ with trivial involutions (cf. Example 3.1.4);
- (2) $\Omega = \mathbb{T}^n$, $\mathbb{k} = \mathbb{C}$, $R = \mathbb{C}[x_1, \dots, x_n]$, $L = \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ with complex conjugation and involution $-\circ$ on L defined as in Example 3.1.5.

Additionally, we fix a filtration $\{R_{\leq d}\}_{d \in \mathbb{N}}$ of R consisting of finite-dimensional vector spaces. Recall that the support of a non-negative or signed measure is defined as follows.

Definition 3.4.1 (cf. [Sch73, Chapter 1.3]). Let μ be a signed measure on Ω . Then

$$\text{supp } \mu := \{ \xi \in \Omega \mid \mu|_U \neq 0 \text{ for all open neighborhoods } U \subseteq \Omega, \xi \in U \}$$

is called *support* of μ , where $\mu|_U$ denotes the restriction of μ to U .

More generally, the same definition applies to distributions. By convention, we consider the support in terms of the standard topology on Ω . The complement of $\text{supp } \mu$ in Ω is the union of all open sets on which μ is constantly zero and is open, so $\text{supp } \mu$ is a closed set. When we consider the support in terms of the Zariski topology, we denote it by $\overline{\text{supp } \mu}$ (as a subset of Ω or $(\mathbb{C}^*)^n$). It is the smallest Zariski-closed set containing $\text{supp } \mu$.

This topic has been studied in various forms, usually in the real affine case with non-negative measures (e. g. [LP15; PPL21]) and an emphasis on finitely-supported measures; see for instance [LR12]. The case of plane algebraic curves has also been investigated in [FAV16], with a focus on the presence of noise. The case of plane trigonometric curves on the torus has been considered in [OJ15; OJ16]. We unify the different noise-free settings in Theorem 3.4.11 and expand the existing results by Theorem 3.4.3, a statement for compactly-supported signed measures.

3.4.1 Signed measures

Here, we consider a signed measure μ on Ω . If $\mathbb{k} = \mathbb{C}$, as in the trigonometric case, then μ is a complex measure. As a consequence of the Riesz representation theorem (see e. g. [Rud87, Theorem 6.19]), these measures can be defined as elements in the continuous dual space of the space $C_c^0(\Omega)$ of compactly-supported continuous functions from Ω to \mathbb{k} . We refer to [Sch73, Chapter 1.2] for an extensive treatment of this topic.

In the trigonometric case, all the moments of μ are defined, as the torus \mathbb{T}^n is compact. In order to speak of moments $\int_{\Omega} x^\alpha d\mu$, $\alpha \in \mathbb{N}^n$, in the affine case, we need to make additional assumptions on the measure μ , since the monomials x^α are not compactly-supported functions on \mathbb{R}^n . Certainly, the moments are defined when the measure μ itself is compactly supported. More generally, all the moments are defined for signed measures with a sufficiently rapid decay toward infinity, such as those that can be written as a product $\mu = g\mu_0$ of a Schwartz function g and a tempered distribution μ_0 (see e. g. [Gra14, Chapter 2] or [Sch73, Chapter 7]), which in particular includes mixtures of Gaussians and mixtures of local mixtures of Gaussians, as considered in Example 2.4.6. In this section, we focus on signed measures with compact support only, as these are determined by their moments.

First, let us take note of the following elementary properties of the support of the product between a measure and a continuous function.

Lemma 3.4.2. *Let μ be a signed measure on Ω and let $f \in C^0(\Omega)$ be a continuous function. Then:*

- (1) $D(f) \cap \text{supp } \mu \subseteq \text{supp}(f\mu)$, where $D(f) \subseteq \Omega$ denotes the set of points in which f does not vanish.

(2) *The measure $f\mu$ is zero if and only if f vanishes on $\text{supp } \mu$.*

Proof. For (1), let $\xi \in \text{supp } \mu$ be any point such that $f(\xi) \neq 0$ and let $U \subseteq \Omega$ be an arbitrary open neighborhood of ξ . We need to show that $f\mu|_U \neq 0$. For this, let $U_0 \subseteq U$ be an open neighborhood of ξ in which f does not have any roots. Since ξ is a support point of μ , there exists a compactly-supported continuous function $\varphi \in C_c^0(U_0)$ such that $\int_{U_0} \varphi d\mu \neq 0$. Then $\psi := \frac{\varphi}{f} \in C_c^0(U_0)$ can be extended trivially to a compactly-supported function $\psi \in C_c^0(U)$ and we have $\int_U \psi d(f\mu) = \int_{U_0} \frac{\varphi}{f} d(f\mu) = \int_{U_0} \varphi d\mu \neq 0$ and thus $f\mu|_U \neq 0$, which proves the statement. For part (2), assume that $f\mu$ is zero. Then $\text{supp}(f\mu) = \emptyset$, so f vanishes on $\text{supp } \mu$ by (1). The converse holds by [Sch73, Chapter 3, Theorem 33, addendum]. \square

For the remainder of this section, we fix a filtration $\{L_{\leq d}\}_{d \in \mathbb{N}}$ of L for which all the components are finite-dimensional vector spaces. In the affine case, we may choose $L_{\leq d} = R_{\leq d}$. Additionally, we denote by B_d^L and B_d^R any bases of the filtered components $L_{\leq d}$ and $R_{\leq d}$, respectively. With this notation, we arrive at the following theorem.

Theorem 3.4.3. *Let μ be a compactly-supported signed measure on Ω , denote by $\mathfrak{a} := I(\text{supp } \mu) \subseteq L$ the vanishing ideal of (the Zariski closure of) its support and let $\sigma: L \rightarrow \mathbb{k}$ be its moment functional. Let $d \in \mathbb{N}$. Then*

$$\mathfrak{a} \cap R_{\leq d} = \ker H_{d',d}$$

holds for all sufficiently large $d' \in \mathbb{N}$, where $H_{d',d} := (\langle w, v \rangle_\sigma)_{w \in B_{d'}^L, v \in B_d^R}$.

It then follows from Hilbert's basis theorem that \mathfrak{a} is generated by $\ker H_{d',d}$ if $d \in \mathbb{N}$ is sufficiently large.

Proof. As the measure μ is compactly supported, all its moments exist. Let $d \in \mathbb{N}$ be arbitrary and observe that

$$\mathfrak{a} \cap R_{\leq d} \subseteq \ker H_{d',d} = \{p \in R_{\leq d} \mid \langle q, p \rangle_\sigma = 0 \text{ for all } q \in L_{\leq d'}\}, \quad (3.4)$$

for all $d' \in \mathbb{N}$. Indeed, if $p \in \mathfrak{a}$, then p vanishes on the support of μ , so $\langle q, p \rangle_\sigma = \int_\Omega q^\circ p d\mu = 0$ for all $q \in L$, by Lemma 3.4.2 (2). More specifically, we have a descending chain

$$R_{\leq d} \supseteq \ker H_{0,d} \supseteq \ker H_{1,d} \supseteq \cdots \supseteq \mathfrak{a} \cap R_{\leq d}$$

which must stabilize, so we can fix a $d' \in \mathbb{N}$ such that

$$\ker H_{d',d} = \ker H_{d'+\delta,d} \quad (3.5)$$

holds for all $\delta \in \mathbb{N}$.

Assume that $\ker H_{d',d} \not\subseteq \mathfrak{a} \cap R_{\leq d}$. Then we can choose a polynomial $p \in \ker H_{d',d}$ with $p \notin \mathfrak{a}$, so p does not vanish everywhere on $\text{supp } \mu$. Hence, by Lemma 3.4.2 (2), the signed measure $\nu := p\mu$ is non-zero, so there exists a compactly-supported continuous function $\varphi \in C_c^0(\Omega)$ such that

$$\int_\Omega \varphi d\nu \neq 0.$$

By the Weierstrass approximation theorem (see [Con90, Chapter 5, Theorem 8.1] for the affine real¹ and [Gra14, Corollary 3.2.2] for the trigonometric version), the function φ can be uniformly approximated by polynomials in L on a compact set containing the support of the measure ν , which implies that not all moments of ν can be zero. Hence, there exists a polynomial $q \in L$ such that $\int_{\Omega} q \, d\nu = \int_{\Omega} qp \, d\mu = \langle q^\circ, p \rangle_\sigma \neq 0$. As $q^\circ \in L_{\leq d'+\delta}$ for some $\delta \in \mathbb{N}$, this implies that $p \notin \ker H_{d'+\delta, d}$, which is a contradiction to (3.5), by the choice of the polynomial p . \square

Remark 3.4.4. In the proof of Theorem 3.4.3, the hypothesis that the support of the signed measure μ is compact does not only guarantee that all its moments exist, but, more importantly, it asserts that the signed measure $\nu = p\mu$ is determined by its moments, so that ν is already zero if all its moments vanish. This does not in general hold for measures that are not compactly supported – not even for rapidly decreasing functions. For instance, let g be a non-zero Schwartz function on \mathbb{R}^n such that all its derivatives vanish at the origin, i. e. $(\partial^\alpha g)(0) = 0$ for all $\alpha \in \mathbb{N}^n$. Then its Fourier transform \hat{g} is a non-zero Schwartz function satisfying

$$(-1)^n (2\pi i)^{|\alpha|} \int_{\mathbb{R}^n} x^\alpha \hat{g}(x) \, dx = (\partial^\alpha g)(0) = 0$$

for all $\alpha \in \mathbb{N}^n$ (cf. [Gra14, Proposition 2.2.11 (10)]), so all the moments of \hat{g} are zero. \diamond

Remark 3.4.5. In the affine case of Theorem 3.4.3, we can choose the filtration of L as $L_{\leq d} = R_{\leq d}$ for all $d \in \mathbb{N}$. Let $H_{d', d}$ be the rectangular moment matrix satisfying the statement of the theorem, so $\mathfrak{a} \cap R_{\leq d} = \ker H_{d', d}$. By Corollary 3.2.5, this equality can only hold when the induced map $\overline{H}_{d', d}$ on the quotient spaces, as in (3.1) of Theorem 3.2.4, is injective. This implies $d' \geq d$, for the affine case. In contrast to the rectangular matrix considered in Theorem 1.3.6, the moment matrix $H_{d', d}$ that we consider here has a different shape, i. e. the dimension of the codomain is not smaller, but larger than or equal to the dimension of the domain of the map. \diamond

By Theorem 3.4.3, we can recover the vanishing ideal of the support from finitely many moments. In particular, this means that the kernel of the non-truncated moment map also yields the vanishing ideal, as the following statement shows.

Corollary 3.4.6. *Under the assumptions of Theorem 3.4.3, we have*

$$\mathfrak{a} \cap R = \ker H,$$

where H denotes the map $H: R \rightarrow \text{Hom}_{\mathbb{k}}^{\text{semi}}(L, \mathbb{k})$, $p \mapsto (q \mapsto \langle q, p \rangle_\sigma)$.

Proof. To see this, first observe that we always have the inclusion $\mathfrak{a} \cap R \subseteq \ker H$, by Lemma 3.4.2(2). On the other hand, if $p \in \ker H$, then $p \in R_{\leq d}$ for some $d \in \mathbb{N}$. In particular, this implies $\langle q, p \rangle_\sigma = 0$ for all $q \in L_{\leq d'} \subseteq L$ and arbitrary $d' \in \mathbb{N}$. Choosing d' as in Theorem 3.4.3, we therefore obtain $p \in \ker H_{d', d} = \mathfrak{a} \cap R_{\leq d}$, so the statement follows. \square

¹This argument would not hold if, in the affine case, we were to work over the field of complex numbers, as the algebra of polynomials on \mathbb{C}^n is not closed under conjugation.

Remark 3.4.7. Theorem 3.4.3 does not quantify what it means for $d' \in \mathbb{N}$ to be large enough for the statement to hold. In general, the choice of d' cannot be made purely based on knowledge of the support or its vanishing ideal, but it must inherently depend on the signed measure itself. Indeed, for arbitrarily large $d, d' \in \mathbb{N}$, one can construct a signed measure with the following properties: its support is compact and Zariski-dense, so its vanishing ideal is zero, and all the low order moments vanish so that the matrix $H_{d',d} = (\langle w, v \rangle_\sigma)_{w \in B_{d'}^L, v \in B_d^R}$ is zero. Hence, the kernel of $H_{d',d}$ is non-zero and thus is not a generating set of the zero ideal, the vanishing ideal of the support. In other words, d' is not large enough for the statement of the theorem to hold. \diamond

Remark 3.4.8. In the trigonometric case, we could also state Theorem 3.4.3 in a more symmetric fashion in terms of a matrix for which both rows and columns are indexed by B_d^L , a basis of the filtered component $L_{\leq d}$. We prefer to index the columns by B_d^R because it allows for a finer filtration, as discussed in Remark 3.3.2. Indexing the rows of the matrix by $B_{d'}^L$ is needed in the proof of Theorem 3.4.3 due to the use of the Weierstrass approximation theorem. This leads to the question whether a statement similar to Theorem 3.4.3 is possible in which rows and columns are indexed by $B_{d'}^R, B_d^R$, i. e. bases of components of the filtration on R instead of L . In general, this is answered negatively by the following example, but a positive answer is possible for non-negative measures, as will be shown in Section 3.4.2. \diamond

Example 3.4.9. We consider the two-dimensional trigonometric case, so let $n = 2$. Let $v_1 := (2, 1), v_2 := (1, 2) \in \mathbb{Z}^2$ and define the functionals

$$\sigma_j: L \longrightarrow \mathbb{C}, \quad x^\alpha \longmapsto \begin{cases} 1 & \text{if } \langle \alpha, v_j \rangle = 0, \\ 0 & \text{otherwise,} \end{cases}$$

for $\alpha \in \mathbb{Z}^2$ and $j = 1, 2$. These are moment functionals of uniform measures supported on the one-dimensional varieties in \mathbb{T}^2 that are defined by the polynomials $x_1 - x_2^2$ and $x_1^2 - x_2$, respectively. Thus, the functional $\sigma := \sigma_1 - \sigma_2$ is a moment functional of a signed measure.

Observe that $\langle x^\alpha, 1 \rangle_\sigma = \sigma(x^{-\alpha}) = 0$ holds for all $\alpha \in \mathbb{N}^2$. This implies that $\langle q, 1 \rangle_\sigma = 0$ for all $q \in R$. Hence, for every choice of d, d' , the polynomial $p := 1$ is contained in the kernel of the moment matrix $H_{d',d} := (\langle w, v \rangle_\sigma)_{w \in B_{d'}^R, v \in B_d^R}$, where B_d^R denotes a basis of $R_{\leq d}$, with respect to any filtration of R . As $p = 1$ does not vanish on any non-empty variety, this shows that the statement of Theorem 3.4.3 does not hold for this matrix $H_{d',d}$ with rows indexed by $B_{d'}^R$ rather than $B_{d'}^L$.

Additionally, this shows that the kernel of the non-truncated map $R \rightarrow \text{Hom}_{\mathbb{k}}^{\text{semi}}(R, \mathbb{k})$, $p \mapsto (q \mapsto \langle q, p \rangle_\sigma)$, is not in general an ideal in R , in the trigonometric case. For instance, we have $\langle x_2^2, x_1 \rangle_\sigma = \sigma(x^{(1,-2)}) \neq 0$, so $x_1 \notin \ker H$, even though $1 \in \ker H$. \diamond

3.4.2 Non-negative measures

Let us consider the case of a (non-negative) measure. The non-negativity is an essential property which allows us to infer stronger results than in the case of signed measures treated in the previous section.

Lemma 3.4.10. *Let μ be a measure on Ω that has finite moments and let $\sigma: L \rightarrow \mathbb{k}$ be its moment functional. If $p \in L$, then*

$$\langle p, p \rangle_\sigma = \int_\Omega |p(x)|^2 d\mu(x) \geq 0.$$

Hence, $\langle -, - \rangle_\sigma$ is positive-semidefinite.

Recall that, in terms of the moment matrix $H_d = (\langle x^\alpha, x^\beta \rangle_\sigma)_{|\alpha|_R, |\beta|_R \leq d}$, we have $\langle p, p \rangle_\sigma = \sigma(p^\circ p) = p^\circ H_d p$, if $p \in R_{\leq d}$ for some $d \in \mathbb{N}$ and p is represented with respect to the monomial basis. Then H_d is a positive-semidefinite matrix.

Proof. In the affine case, p is a real polynomial and we have $\sigma(pp) = \int_{\mathbb{R}^n} p(x)^2 d\mu(x) \geq 0$. In the trigonometric case, p is a trigonometric polynomial, so that $\overline{p^\circ(\xi)} = p(\xi)$ for all $\xi \in \mathbb{T}^n$, as mentioned in Example 3.1.5. Therefore, $\sigma(p^\circ p) = \int_{\mathbb{T}^n} \overline{p(x)} p(x) d\mu(x) = \int_{\mathbb{T}^n} |p(x)|^2 d\mu(x) \geq 0$. \square

Note that a converse of Lemma 3.4.10 does not generally hold, i.e. if $\sigma: L \rightarrow \mathbb{k}$ is any functional such that $\langle -, - \rangle_\sigma$ is positive-semidefinite, we cannot conclude that there exists a measure that has σ as moment functional, since not every non-negative multivariate polynomial is a sum of squares. See [Sch17] for an extensive treatment of questions pertaining to the existence of representing measures.

The following theorem is a stronger version of Theorem 3.4.3 for the case of (non-negative) measures. If $W = R_{\leq d}$ is a component of the total degree filtration in the affine case, this statement can also be obtained by combining [LR12, Theorem 2.10] and [PPL21, Lemma 5].

Theorem 3.4.11. *Let μ be a measure on Ω with finite moments, let $\mathfrak{a} := I(\text{supp } \mu) \subseteq L$ be the vanishing ideal of (the Zariski closure of) its support and let $\sigma: L \rightarrow \mathbb{k}$ be its moment functional. Let $W \subseteq R$ be a \mathbb{k} -vector subspace. Then $\langle -, - \rangle_\sigma$ induces a positive-definite form on $W/(\mathfrak{a} \cap W)$.*

In particular, if W is finite-dimensional and B is a basis of W , let $H := (\langle w, v \rangle_\sigma)_{w, v \in B}$. Then

$$\mathfrak{a} \cap W = \ker H.$$

Furthermore, H is non-singular if and only if the elements of B are linearly independent modulo $\mathfrak{a} \cap W$.

For the statement, only finiteness of the moments that occur in H is needed, so σ must be defined on the subspace $W^\circ \cdot W \subseteq L$.

Proof. First observe that $\langle -, - \rangle_\sigma$ is positive-semidefinite by Lemma 3.4.10. Then it follows by Lemmas 3.1.16 and 3.1.17 that it is enough to show that $\langle -, - \rangle_\sigma$ induces a non-degenerate form on the space $W/(\mathfrak{a} \cap W)$ in order to conclude that $\mathfrak{a} \cap W = \ker H$, as H agrees with the map from Lemma 3.1.17 (3).

From this, the addendum readily follows. If H is non-singular, we have $\mathfrak{a} \cap W = \ker H = 0$, so the elements of B are linearly independent modulo $\mathfrak{a} \cap W$. If H is singular, we find a non-trivial linear combination $q = \sum_{w \in B} q_w w \neq 0$, $q_w \in \mathbb{k}$, with $q \in \ker H = \mathfrak{a} \cap W$, so $q \equiv 0 \pmod{\mathfrak{a} \cap W}$.

By Lemma 3.1.11, $\langle -, - \rangle_\sigma$ induces a form on $W/(\mathfrak{a} \cap W)$. It remains to show that it is non-degenerate. For this, assume that $p \in W$ is a polynomial such that $\langle p, p \rangle_\sigma = 0$. Since $|p|^2 \geq 0$ on Ω , it follows from [Sch17, Proposition 1.23] that $|p|^2$ vanishes on $\text{supp } \mu$ and thus $p \in \mathfrak{a}$. \square

In particular, Theorem 3.4.11 holds with $W = R_{\leq d}$ for any $d \in \mathbb{N}$, so that

$$\mathfrak{a} \cap R_{\leq d} = \ker H.$$

Again, by Hilbert's basis theorem, the ideal \mathfrak{a} is generated by $\mathfrak{a} \cap R_{\leq d}$ if d is sufficiently large. Hence, for such a number d , we can fully recover the ideal \mathfrak{a} from finitely many moments, namely from $\ker H$, which is the statement of [LR12, Theorem 2.10].

3.4.3 Examples

We give a few examples of signed measures that illustrate that the assumption of non-negativity is crucial for Theorem 3.4.11.

Example 3.4.12. Let μ be a signed measure supported on the real interval $[-1, 1] \subseteq \mathbb{R}$ with density x and denote its moment functional by $\sigma: R := \mathbb{R}[x] \rightarrow \mathbb{R}$, so that $\sigma(p) = \int_{-1}^1 p(x)x dx$ for $p \in R$. In particular, this means that $\langle -, - \rangle_\sigma$ is not positive-semidefinite. One checks that, due to symmetry, the even moments $\sigma(x^{2\alpha}) = 0$ vanish for $\alpha \in \mathbb{N}$ and thus $\det(\sigma(x^{2\alpha+2\beta}))_{0 \leq \alpha, \beta \leq d} = 0$ for all $d \in \mathbb{N}$. Then it follows that $\det(\sigma(x^{\alpha+\beta}))_{0 \leq \alpha, \beta \leq d} = 0$ if d is even, for example using the Leibniz formula or by a suitable permutation of rows and columns.

This means that, for every even d , we find some non-zero polynomial in $R_{\leq d}$ that lies in the kernel of the moment matrix $(\sigma(x^{\alpha+\beta}))_{0 \leq \alpha, \beta \leq d}$, even though the variety corresponding to the Zariski closure of the support of the signed measure μ is the entire line \mathbb{R} , which is defined by the zero-ideal in R , and despite the fact that the monomials are linearly independent modulo the zero-ideal. Hence, the statement of Theorem 3.4.11 cannot hold. However, note that, in this example, the non-truncated Hankel operator is injective nevertheless, as stated in Corollary 3.4.6. \diamond

In the affine setting with $L_{\leq d} = R_{\leq d}$, $d \in \mathbb{N}$, and for a finitely-supported signed measure, it follows from Lemma 1.3.3 that the statement of Theorem 3.4.3 holds with $d' := d$, as

long as $d \in \mathbb{N}$ is sufficiently large. The following example shows that this can fail for small d .

Example 3.4.13. Let $R = \mathbb{k}[x]$ be the univariate polynomial ring and let $\mathfrak{a} = \mathfrak{m}_{\xi_1} \cap \mathfrak{m}_{\xi_2}$ with two distinct points $\xi_1, \xi_2 \in \mathbb{k}$. We consider the map $\sigma = \text{ev}_{\xi_1} - \text{ev}_{\xi_2}$. Denote by $H_{d',d}$ the corresponding Hankel matrix, for $d, d' \in \mathbb{N}$. By (3.4), we have $\mathfrak{a} \cap R_{\leq d} \subseteq \ker H_{d',d}$, but equality does not hold for small d .

For instance, if $d' = d = 0$, we have

$$\mathfrak{a} \cap R_{\leq d} = 0 \subsetneq \ker H_{0,0} = \ker(0).$$

However, if d is sufficiently large, namely $d \geq 2$, and if $d' \geq d$, we have $\mathfrak{a} \cap R_{\leq d} = \ker H_{d',d}$ by Lemma 1.3.3, regardless of the choice of d' . \diamond

In contrast, we have seen in Example 3.4.9 that a similar statement is not possible for infinitely-supported signed measures. More precisely, it is an example in which one has $\mathfrak{a} \cap R_{\leq d} \neq \ker H_{d,d}$ for all $d \in \mathbb{N}$, since $1 \in \ker H_{d,d}$, but $1 \notin \mathfrak{a}$. For a non-negative measure, this would not be possible due to Theorem 3.4.11.

3.5 From moments to approximations of measures

As we have seen in Section 3.4, it is possible to recover the Zariski closure of the support of a non-negative or a compactly-supported signed measure, by an algebraic computation involving finitely many moments. The next step consists of recovering the remaining unknown data that define the measure, such as information about a density of the measure. This is still a broad problem. One way to address this is to consider measures with densities defined by finitely many parameters, such as piecewise-constant or -polynomial functions. This is the approach of e.g. [PT14; OJ15; OJ16; FAV16].

Here, we follow a different approach. Rather than attempting to find the defining parameters exactly, we construct several functions from the moments that approximate certain aspects of the measure. On the one hand, this is related to MUSIC [Sch86], a recovery method for finitely-supported measures. On the other hand, it is connected to Christoffel functions (see e.g. [Nev86] for an overview), a topic with a long history that mainly pertains to measures that are fully supported on the whole space. Recently, Christoffel functions have been studied in the context of measures supported on algebraic varieties, as well (cf. [PPL21; Mar+21]). In the case of finitely-supported measures, both of these areas are connected to beamforming techniques in signal processing; see e.g. [KV96] and the references therein.

In this section, we focus on the trigonometric setting on the torus. Here, we identify the complex torus with the unit cube $\mathbf{T}^n := [0, 1)^n$, $n \in \mathbb{N}$. If we consider the torus embedded into \mathbb{C}^n , we denote it by $\Omega := \mathbb{T}^n = \iota(\mathbf{T}^n) \subseteq \mathbb{C}^n$, using the parametrization

$$\iota: \mathbf{T}^n = [0, 1)^n \longrightarrow \mathbb{T}^n \subseteq \mathbb{C}^n, \quad (t_1, \dots, t_n) \longmapsto (e^{2\pi i t_1}, \dots, e^{2\pi i t_n}).$$

We use variables t, ϑ to refer to points in \mathbf{T}^n and $x = \iota(t)$, $\xi = \iota(\vartheta)$ to denote the corresponding points in Ω . The torus \mathbf{T}^n is endowed with the topology that is induced by Ω .

3.5.1 Definitions and first properties

As before, let $R = \mathbb{C}[x_1, \dots, x_n]$ denote the polynomial ring and fix the filtration $\{R_{\leq d}\}_{d \in \mathbb{N}}$ defined by max-degree. Moreover, denote by $L = \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ the Laurent polynomial ring endowed with involution $-\circ$ as defined in Example 3.1.5. Thus, a trigonometric polynomial $\sum_{\alpha} p_{\alpha} e^{2\pi i \langle \alpha, t \rangle} = \sum_{\alpha} p_{\alpha} \iota(t)^{\alpha}$ with finitely many non-zero coefficients $p_{\alpha} \in \mathbb{C}$, $\alpha \in \mathbb{Z}^n$, is the same as the composition of the embedding ι with the Laurent polynomial $\sum_{\alpha} p_{\alpha} x^{\alpha} \in L$. Also recall that $p^{\circ}(\xi) = \overline{p(\xi)}$ holds for all $p \in L$ and $\xi \in \Omega$.

Definition 3.5.1. For every $d \in \mathbb{N}$, we fix the basis B_d of $R_{\leq d}$ defined by

$$B_d := \left\{ (d+1)^{-\frac{n}{2}} x^{\alpha} \mid \alpha \in \mathbb{N}^n, |\alpha|_{\infty} \leq d \right\}.$$

This basis satisfies an addition theorem of the form

$$\sum_{w \in B_d} w^{\circ} w = (d+1)^{-n} \sum_{|\alpha|_{\infty} \leq d} x^{-\alpha} x^{\alpha} = 1. \quad (3.6)$$

We use the notation $W_d := R_{\leq d}$ to emphasize that we work with the basis B_d , rather than with the monomial basis of $R_{\leq d}$.

Definition 3.5.2. For $\xi \in \Omega$, define the vector

$$e_{d\xi} := \sum_{w \in B_d} \overline{w(\xi)} w \in W_d,$$

where $w(\xi) \in \mathbb{C}$ is the evaluation of w at ξ . We denote the sesquilinear inner product on the vector space W_d with respect to the basis B_d by

$$\langle -, - \rangle_{B_d} : W_d \times W_d \longrightarrow \mathbb{C}, \quad \left(\sum_{w \in B_d} q_w w, \sum_{w \in B_d} p_w w \right) \longmapsto \sum_{w \in B_d} \overline{q_w} p_w.$$

The vector $e_{d\xi} \in W_d$ is dual to the evaluation functional ev_{ξ} in the sense that, for any $p = \sum_{w \in B_d} p_w w \in W_d$, we have

$$\langle e_{d\xi}, p \rangle_{B_d} = \sum_{w \in B_d} \overline{w(\xi)} p_w = p(\xi)$$

and thus

$$\langle e_{d\xi}, - \rangle_{B_d} = \text{ev}_{\xi}|_{W_d} \in W_d^*.$$

This justifies the notation

$$e_{d\xi}^* p = p(\xi),$$

for any $p \in W_d$.

If d is even and if $\xi = \iota(\vartheta)$ for some $\vartheta \in \mathbf{T}^n$, we have

$$e_{d\xi}(\iota(t)) = \sum_{\substack{\alpha \in \mathbb{N}^n \\ |\alpha|_\infty \leq d}} \frac{e^{2\pi i \langle \alpha, t - \vartheta \rangle}}{(d+1)^n} = \prod_{i=1}^n \frac{e^{\pi i d (t_i - \vartheta_i)}}{d+1} \cdot \frac{\sin((d+1)\pi(t_i - \vartheta_i))}{\sin(\pi(t_i - \vartheta_i))},$$

so the trigonometric polynomial $e_{d\xi} \circ \iota$ is the product of a complex exponential and a multivariate *Dirichlet kernel* centered at the point ϑ (cf. [Gra14, Section 3.1.3]).

Definition 3.5.3. If μ is a measure on \mathbf{T}^n , we define the moment functional $\sigma: L \rightarrow \mathbb{C}$ associated to the measure μ as follows. The α -th *trigonometric moment* is given by the $(-\alpha)$ -th Fourier coefficient of μ (cf. [Gra14, Section 3.1]), so

$$\sigma(x^\alpha) := \hat{\mu}(-\alpha) = \int_{\mathbf{T}^n} e^{-2\pi i \langle -\alpha, t \rangle} d\mu(t),$$

for $\alpha \in \mathbb{Z}^n$.

Note that a measure on \mathbf{T}^n always has finite moments due to the compactness of the torus. With the above convention, we have for instance for a Dirac measure δ_ϑ with $\vartheta \in \mathbf{T}^n$ and $\xi = \iota(\vartheta) \in \Omega$ that the moments are of the form

$$\sigma(x^\alpha) = \int_{\mathbf{T}^n} e^{-2\pi i \langle -\alpha, t \rangle} d\delta_\vartheta(t) = e^{2\pi i \langle \alpha, \vartheta \rangle} = \xi^\alpha,$$

as usual.

Associated to σ , we have the sesquilinear form $\langle -, - \rangle_\sigma$ on L , which is positive-semidefinite by Lemma 3.4.10. For $d \in \mathbb{N}$, we denote the moment matrix by

$$H_d := (\langle w, v \rangle_\sigma)_{w, v \in B_d} = (\sigma(w^\circ v))_{w, v \in B_d}.$$

Let $H_d = U\Sigma V^*$ be the singular value decomposition. As H_d is a positive-semidefinite matrix, we can assume that $V = U$, so the singular value decomposition is of the form $H_d = U\Sigma U^*$, where U is a unitary matrix, i. e. its columns are orthonormal with respect to $\langle -, - \rangle_{B_d}$ on W_d . If H_d is real, then also U is a real matrix. Denote the ordered non-zero singular values of H_d by $\varsigma_1 \geq \dots \geq \varsigma_r > 0$, where $r = \text{rk } H_d$, and denote the singular vectors, the columns of U , by u_1, \dots, u_N , where $N := (d+1)^n$ is the cardinality of B_d . Moreover, denote by U_1 and U_0 the truncation of U to the columns u_1, \dots, u_r and u_{r+1}, \dots, u_N , respectively, so in particular the column space of U_0 is $\ker H_d$. We view the singular vectors $u_j = \sum_{w \in B_d} u_{j,w} w \in W_d$ as elements in the vector space W_d and we have $e_{d\xi}^* u_j = u_j(\xi)$ as well as $u_j^* e_{d\xi} = \overline{u_j(\xi)}$ for all $1 \leq j \leq N$.

Definition 3.5.4. With the above notation, we define the following functions associated to the moment functional σ of a (non-negative) measure on \mathbf{T}^n :

$$P_d: \Omega \longrightarrow \mathbb{R}_{\geq 0}, \quad \xi \longmapsto \langle e_{d\xi}, e_{d\xi} \rangle_\sigma = e_{d\xi}^* H_d e_{d\xi},$$

$$\begin{aligned}
 P_{d,1}: \Omega &\longrightarrow \mathbb{R}_{\geq 0}, & \xi &\longmapsto \sum_{j=1}^r |u_j(\xi)|^2 = e_{d\xi}^* U_1 U_1^* e_{d\xi}, \\
 Q_{d,0}: \Omega &\longrightarrow \mathbb{R}_{>0} \cup \{\infty\}, & \xi &\longmapsto \begin{cases} \frac{1}{\sum_{j=r+1}^N |u_j(\xi)|^2} & \text{if } \sum_{j=r+1}^N |u_j(\xi)|^2 \neq 0, \\ \infty & \text{otherwise,} \end{cases} \\
 Q_{d,\varepsilon}: \Omega &\longrightarrow \mathbb{R}_{>0}, & \xi &\longmapsto \frac{1}{e_{d\xi}^* \left(H_d^\dagger + \frac{1}{\varepsilon} U_0 U_0^* \right) e_{d\xi}} \\
 & & &= \frac{1}{\sum_{j=1}^r \frac{1}{\varsigma_j} |u_j(\xi)|^2 + \frac{1}{\varepsilon} \sum_{j=r+1}^N |u_j(\xi)|^2},
 \end{aligned}$$

where $\varepsilon > 0$. Note that these definitions depend on our choice of the basis B_d . See Figures 3.1 and 3.2 for a visualization of these functions.

Observe that $P_d \circ \iota$ and $P_{d,1} \circ \iota$ are trigonometric polynomials on \mathbf{T}^n , but $Q_{d,0} \circ \iota$ and $Q_{d,\varepsilon} \circ \iota$ are not. Moreover, note that $Q_{d,0}(\xi) = \infty$ for all $\xi \in \Omega$ if H_d is non-singular. This happens if the support of the measure is not contained in an algebraic variety or, more specifically, if there exist no polynomials in W_d that describe algebraic relations between the low-order moments.

The equality in the definition of $Q_{d,\varepsilon}$ follows from the fact that the (Moore–Penrose) pseudo-inverse of H_d is of the form

$$H_d^\dagger = U_1 \operatorname{diag}(\varsigma_1^{-1}, \dots, \varsigma_r^{-1}) U_1^*$$

(cf. [HJ13, Problem 7.3.P7]). In order to see that the denominator of $Q_{d,\varepsilon}$ is always non-zero, note that, due to the addition theorem (3.6), we have

$$\begin{aligned}
 \sum_{j=1}^N |u_j(\xi)|^2 &= e_{d\xi}^* U U^* e_{d\xi} = e_{d\xi}^* e_{d\xi} = \langle e_{d\xi}, e_{d\xi} \rangle_{B_d} \\
 &= \sum_{w \in B_d} w(\xi) \overline{w(\xi)} = \left(\sum_{w \in B_d} w w^\circ \right) (\xi) = 1,
 \end{aligned} \tag{3.7}$$

for any choice of $\xi \in \Omega$ and $d \in \mathbb{N}$, since U is unitary. Hence, there is at least one j for which $|u_j(\xi)|^2 > 0$. Finally, observe that $Q_{d,0}$ and $Q_{d,\varepsilon}$ are well-defined, as they do not depend on the concrete choice of the singular vectors u_j .

Remark 3.5.5. If the moment matrix H_d is regular for all $d \in \mathbb{N}$, then $Q_{d,\varepsilon}$ does not depend on ε and we simply denote it by Q_d . In this case, Q_d is equal to the variational form of the *Christoffel function* associated to σ , that is

$$Q_d(\xi) = \min \left\{ \frac{\langle p, p \rangle_\sigma}{|p(\xi)|^2} \mid p \in W_d, p(\xi) \neq 0 \right\} \tag{3.8}$$

for all $\xi \in \Omega$.

If the matrix H_d is singular for some $d \in \mathbb{N}$, the support of the measure is contained in an algebraic variety of dimension less than n , by Theorem 3.4.11. In this case, different generalizations of the Christoffel function exist. See for instance [Mar+21] for an overview. Here, we call the function $Q_{d,\varepsilon}$ *regularized Christoffel function*. The regularization we choose here is convenient from a noise-free theoretical perspective, as it allows for the precise statement of Proposition 3.5.11, for instance. For numerical computations, we determine the (numerical) rank of the moment matrix H_d as the largest integer r such that $\varsigma_r \geq \varepsilon$, in which case the regularization scheme is the same as the *spectral cut-off* regularization mentioned in [Mar+21]. \diamond

Remark 3.5.6. In case of a finitely-supported measure, $Q_{d,0}$ is the function of central interest in the parameter recovery method MUSIC [Sch86], in a noise-free setting. In MUSIC, the square root of $Q_{d,0}$ is also referred to as *imaging function*, while the square root of $1 - P_{d,1}$ corresponds to the *noise-space correlation function* (cf. [PPST18, Chapter 10.2.1]). See also [KV96] for a broader overview. As such, our definitions of these functions may be interpreted as generalizations thereof to the case of infinitely-supported measures. \diamond

In the remainder of this section, we list several interpolating properties and qualitative convergence results for the functions defined above, some of which are unique to the zero-dimensional case, that is, to finitely-supported measures on the torus. The subject of quantifying these convergence properties in terms of explicit convergence rates is left for further study.

We start with the following statement about $P_{d,1}$, which establishes an important link between the support of the measure and $P_{d,1}$, for large $d \in \mathbb{N}$.

Lemma 3.5.7. *Let μ be a measure on \mathbf{T}^n and let $\mathfrak{a} := \mathbb{I}(\iota(\text{supp } \mu)) \subseteq L$ be the vanishing ideal of the support. Then, for any $d \in \mathbb{N}$:*

- (1) $0 \leq P_{d,1}(\xi) \leq 1$ for all $\xi \in \Omega$.
- (2) $P_{d,1}(\xi) = 1$ for all $\xi \in \overline{\iota(\text{supp } \mu)} \subseteq \Omega$.
- (3) *If \mathfrak{a} is generated by $\mathfrak{a} \cap R_{\leq d}$ and if $\xi \in \Omega$, then $P_{d,1}(\xi) = 1$ if and only if $\xi \in \overline{\iota(\text{supp } \mu)} \subseteq \Omega$.*

In other words, $P_{d,1}$ interpolates 1 on the smallest variety containing $\text{supp } \mu$ if d is sufficiently large, as illustrated in Figure 3.2. Here $\overline{\iota(\text{supp } \mu)}$ denotes the Zariski closure in Ω .

Proof. For all $\xi \in \Omega$, we have

$$0 \leq P_{d,1}(\xi) = \sum_{j=1}^r |u_j(\xi)|^2 \leq \sum_{j=1}^N |u_j(\xi)|^2 = 1,$$

where the last equality is due to (3.7).

By Theorem 3.4.11, we have $\mathfrak{a} \cap R_{\leq d} = \ker H_d = \langle u_{r+1}, \dots, u_N \rangle$. Thus, if $\xi \in \overline{\iota(\text{supp } \mu)}$, then $u_j(\xi) = 0$ for all $r < j \leq N$ and therefore $\sum_{j=r+1}^N |u_j(\xi)|^2 = 0$, which is equivalent to $P_{d,1}(\xi) = 1$.

Conversely, if $P_{d,1}(\xi) = 1$ and thus $\sum_{j=r+1}^N |u_j(\xi)|^2 = 0$, we must have $u_j(\xi) = 0$ for all $r < j \leq N$. Since these polynomials span $\mathfrak{a} \cap R_{\leq d}$, it follows that $\xi \in \overline{\iota(\text{supp } \mu)}$, as long as \mathfrak{a} is generated by $\mathfrak{a} \cap R_{\leq d}$. \square

Remark 3.5.8. With Lemma 3.5.7, we can describe $\overline{\iota(\text{supp } \mu)}$ as the zero set in Ω of a single polynomial, namely $1 - P_{d,1}$, for sufficiently large $d \in \mathbb{N}$. This is in contrast to the multivariate Prony method, such as Theorem 1.3.1, which uses multiple polynomials to describe the support of a (finitely-supported) measure as a zero set.

Although we have phrased Lemma 3.5.7 only for positive measures, it can be transferred to finitely-supported signed complex measures on Ω , as well. In this case, the moment matrix may not be positive-semidefinite, so, if $H_d = U_1 \text{diag}(\varsigma_1, \dots, \varsigma_r) V_1^*$ denotes the truncated singular value decomposition, one defines $P_{d,1}(\xi) = e_{d\xi}^* V_1 V_1^* e_{d\xi} \in \mathbb{R}_{\geq 0}$, for $\xi \in \Omega$, purely in terms of the right singular vectors V_1 , disregarding U_1 , as the right kernel of H_d is orthogonal to V_1 rather than U_1 . Note that statement (2) of Lemma 3.5.7 only holds for sufficiently large d , in this case, namely under the assumptions of (3), since only then the equality $\mathfrak{a} \cap R_{\leq d} = \ker H_d$ holds which is used in the proof. As we have seen in Examples 3.4.9 and 3.4.12, this does not directly translate to the case of signed measures that are not finitely-supported. Though, it is possible to consider the right singular vectors of suitable rectangular moment matrices as in Theorem 3.4.3. \diamond

Also note that, if $1 - P_{d,1}$ is extended to a Laurent polynomial on the entire algebraic torus $(\mathbb{C}^*)^n$, it may have additional zeros outside of Ω . Moreover, we have the following characterization, which is of a similar form as the variational formulation (3.8) of the Christoffel function.

Lemma 3.5.9. *Let μ be a measure on \mathbf{T}^n and let $\mathfrak{a} := \mathbb{I}(\iota(\text{supp } \mu)) \subseteq L$ be the vanishing ideal of the support. Then*

$$1 - P_{d,1}(\xi) = \max \left\{ \frac{|q(\xi)|^2}{\|q\|_2^2} \mid q \in \mathfrak{a} \cap W_d \setminus \{0\} \right\} \cup \{0\},$$

for all $d \in \mathbb{N}$ and $\xi \in \Omega$. The maximum is attained at $q = U_0 U_0^* e_{d\xi}$.

Proof. If $\mathfrak{a} \cap W_d = 0$, then the moment matrix H_d is non-singular by Theorem 3.4.11 and thus $P_{d,1}(\xi) = 1$ by (3.7), so we can assume that $\mathfrak{a} \cap W_d \neq 0$. The columns of the matrix U_0 form an orthonormal basis of $\mathfrak{a} \cap W_d$, so the polynomial $q := U_0 U_0^* e_{d\xi} \in \mathfrak{a} \cap W_d$ is the orthogonal projection of $e_{d\xi} \in W_d$ onto the subspace $\mathfrak{a} \cap W_d$. Similarly, if $p \in \mathfrak{a} \cap W_d$ is any polynomial, we have

$$q^* p = e_{d\xi}^* U_0 U_0^* p = e_{d\xi}^* p = p(\xi).$$

In particular, note that

$$\|q\|_2^2 = q^* q = q(\xi) = e_{d\xi}^* U_0 U_0^* e_{d\xi} = 1 - P_{d,1}(\xi). \quad (3.9)$$

Therefore, by the Cauchy–Schwarz inequality, it follows that

$$|p(\xi)|^2 = |q^*p|^2 \leq \|q\|_2^2 \cdot \|p\|_2^2 = (1 - P_{d,1}(\xi))\|p\|_2^2.$$

Hence, we have

$$1 - P_{d,1}(\xi) \geq \max \left\{ \frac{|p(\xi)|^2}{\|p\|_2^2} \mid p \in \mathfrak{a} \cap W_d \setminus \{0\} \right\} \geq \frac{|q(\xi)|^2}{\|q\|_2^2} = 1 - P_{d,1}(\xi),$$

if $q \neq 0$. The first inequality also holds when $q = 0$, in which case the result follows due to (3.9). \square

Lemma 3.5.10. *Let μ be a measure on \mathbf{T}^n and let $\mathfrak{a} := \mathbf{I}(\iota(\text{supp } \mu)) \subseteq L$. Then, for all $d \in \mathbb{N}$:*

- (1) $\frac{1}{\varepsilon}Q_{d,\varepsilon} \leq Q_{d,0}$ for all $\varepsilon > 0$.
- (2) $\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon}Q_{d,\varepsilon}(\xi) = Q_{d,0}(\xi)$ for all $\xi \in \Omega$.
- (3) $Q_{d,0}(\xi) = \infty$ for all $\xi \in \overline{\iota(\text{supp } \mu)} \subseteq \Omega$.
- (4) If \mathfrak{a} is generated by $\mathfrak{a} \cap R_{\leq d}$, then $Q_{d,0}(\xi) < \infty$ for $\xi \in \Omega \setminus \overline{\iota(\text{supp } \mu)}$.

Here $\overline{\iota(\text{supp } \mu)}$ denotes the Zariski closure of $\iota(\text{supp } \mu)$ in Ω .

Proof. Part (1) follows immediately from the definitions. Moreover, if $\sum_{j=r+1}^N |u_j(\xi)|^2 = 0$, then

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon}Q_{d,\varepsilon}(\xi) = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon \sum_{j=1}^r \frac{1}{\zeta_j} |u_j(\xi)|^2} = \infty = Q_{d,0}(\xi).$$

Otherwise, we have

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon}Q_{d,\varepsilon}(\xi) = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon \sum_{j=1}^r \frac{1}{\zeta_j} |u_j(\xi)|^2 + \sum_{j=r+1}^N |u_j(\xi)|^2} = Q_{d,0}(\xi) < \infty,$$

which settles (2). Furthermore, it follows from (3.7) that we have

$$Q_{d,0}(\xi) = \frac{1}{1 - P_{d,1}(\xi)} \tag{3.10}$$

for all $\xi \in \Omega$, with the convention that $\frac{1}{0} = \infty$. Hence, the remaining properties are an immediate consequence of Lemma 3.5.7. \square

We can now show that $Q_{d,\varepsilon}$ interpolates the weights of a finitely-supported measure on Ω if $d \in \mathbb{N}$ is sufficiently large. A similar statement for the real affine setting can be found in [Sch17, Theorem 18.42].

Proposition 3.5.11. *Let $\mu = \sum_{j=1}^r \lambda_j \delta_{\xi_j}$ be a (non-negative) measure supported at finitely many distinct points $\xi_1, \dots, \xi_r \in \Omega$, where $\lambda_1, \dots, \lambda_r \in \mathbb{R}_{>0}$. If $d \in \mathbb{N}$ is large enough such that $\text{rk } H_d = r$, then*

$$Q_{d,\varepsilon}(\xi_j) = \lambda_j$$

for all $1 \leq j \leq r$ and $\varepsilon > 0$.

Proof. By Theorem 3.2.4 and Example 3.2.7, we obtain a factorization $H_d = A^* \Lambda A$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_r)$ and where A is the natural quotient map from W_d to $S := \bigoplus_{j=1}^r R/\mathfrak{m}_{\xi_j}$ if d is large enough. Note that $A = (w(\xi_j))_{1 \leq j \leq r, w \in B_d} = (e_{d\xi_j})_{1 \leq j \leq r}^*$, which, up to a scalar factor, is a Vandermonde matrix.

As the matrix A must have full row rank r , the (Moore–Penrose) pseudo-inverse A^\dagger is a right-inverse of A . Thus, denoting the standard basis of S by e_1, \dots, e_r , then, for every $1 \leq j \leq r$, the element $\ell_j := A^\dagger e_j \in W_d$ is a Lagrange polynomial (cf. [Isk18, Chapter 8.1]), as it must satisfy $\ell_j(\xi_k) = \delta_{jk}$ for all $1 \leq k \leq r$, since $A\ell_j = AA^\dagger e_j = e_j$. In particular, for all $\xi \in \Omega$, we can write

$$e_{d\xi}^* A^\dagger = (\ell_1(\xi), \dots, \ell_r(\xi)).$$

As A has full row rank, the pseudo-inverse of H_d is $H_d^\dagger = A^\dagger \Lambda^{-1} (A^*)^\dagger$ (cf. [Bjö96, Theorem 1.2.13]). Therefore, it follows that

$$Q_{d,\varepsilon}(\xi) = \frac{1}{e_{d\xi}^* \left(A^\dagger \Lambda^{-1} (A^*)^\dagger + \frac{1}{\varepsilon} U_0 U_0^* \right) e_{d\xi}} = \frac{1}{\sum_{j=1}^r \frac{1}{\lambda_j} |\ell_j(\xi)|^2 + \frac{1}{\varepsilon} \sum_{j=r+1}^N |u_j(\xi)|^2},$$

for all $\xi \in \Omega$. Then, due to Theorem 3.4.11, the statement follows from the observation that, for $r < j \leq N$ and $\xi \in \{\xi_1, \dots, \xi_r\}$, we have $u_j(\xi) = 0$ since $u_j \in \ker H_d$. \square

Remark 3.5.12. More generally, one can define the functions in Definition 3.5.4 on other measurable spaces Ω for which an addition theorem as in (3.6) is satisfied. For instance, this is possible on the real sphere $\mathbb{S}^{n-1} \subseteq \mathbb{R}^n$ as well as the rotation group $\text{SO}(3)$. The proofs of Lemmas 3.5.7, 3.5.9 and 3.5.10 and Proposition 3.5.11 only depend on such an addition theorem as well as the property $p^\circ(\xi) = \overline{p(\xi)}$ for $p \in L$ and $\xi \in \Omega$, so these statements can be transferred, as well. See for instance [KMvdO19].

Another possible generalization consists of defining the functions relatively to a suitable reference measure, which is useful for spaces that are not compact. This approach is pursued in [PPL21]. In our case, the reference measure is the n -dimensional Lebesgue measure on the torus, up to normalization. \diamond

3.5.2 Convergence results

In the following, we show qualitative convergence properties of the functions from Definition 3.5.4, as the degree d is increased. These statements depend on properties that are inherent to the torus. First, observe that, for every polynomial $q \in \sum_{|\alpha|_\infty \leq d} q_\alpha \frac{x^\alpha}{\sqrt{N}} \in W_d$, $d \in \mathbb{N}$, written in the basis B_d , *Plancherel's identity* (cf. [Gra14, Proposition 3.2.7]) reads as

$$\int_{\Omega} |q(x)|^2 dx = \sum_{|\alpha|_\infty, |\beta|_\infty \leq d} \frac{q_\alpha \overline{q_\beta}}{N} \int_{\Omega} x^{\alpha-\beta} dx = \frac{\|q\|_2^2}{N} \quad (3.11)$$

where $N := (d+1)^n$ denotes the dimension of W_d . We can now prove the following theorem about the pointwise limiting behavior of the functions $P_{d,1}$, $Q_{d,0}$ and $Q_{d,\varepsilon}$.

Theorem 3.5.13. *Let μ be a measure on \mathbf{T}^n and let $\varepsilon > 0$. Then*

$$\lim_{d \rightarrow \infty} P_{d,1}(\xi) = 0, \quad \lim_{d \rightarrow \infty} Q_{d,0}(\xi) = 1, \quad \limsup_{d \rightarrow \infty} Q_{d,\varepsilon}(\xi) \leq \varepsilon$$

for all $\xi \in \Omega \setminus \overline{\iota(\text{supp } \mu)}$.

In particular, this means that the function $P_{d,1}$ converges pointwisely to the indicator function of $\overline{\iota(\text{supp } \mu)}$, by Lemma 3.5.7 (2), where $\overline{\iota(\text{supp } \mu)}$ denotes the Zariski closure in Ω .

Proof. If $\overline{\iota(\text{supp } \mu)} = \Omega$, there is nothing to show. Otherwise, $\iota(\text{supp } \mu)$ is contained in a variety of dimension less than n , as Ω is Zariski-dense in \mathbb{C}^n , which is generated by a non-zero ideal $\mathfrak{a} \subseteq L$. Thus, for every sufficiently large $d \in \mathbb{N}$, we have $\mathfrak{a} \cap W_d \neq 0$. In order to prove the statement about $P_{d,1}$, it is therefore enough to show that the maximum of $\left\{ \frac{|q(\xi)|^2}{\|q\|_2^2} \mid q \in \mathfrak{a} \cap W_d \setminus \{0\} \right\}$ converges to 1, for $d \rightarrow \infty$, due to Lemma 3.5.9.

Define the locally compact space $X := \Omega \setminus \overline{\iota(\text{supp } \mu)}$ and denote by $C_0^0(X)$ the complex-valued continuous functions vanishing at infinity, which we identify with the continuous functions on Ω vanishing on the boundary $\Omega \setminus X = \overline{\iota(\text{supp } \mu)}$. Let $f \in C_0^0(X)$ be any function satisfying $\sup_X |f| = |f(\xi)| = 1$. Note that we can view \mathfrak{a} as a (non-unital) complex subalgebra of $C_0^0(X)$ by restricting the Laurent polynomials to $X \subseteq \Omega$. It is closed under complex conjugation, as $\mathfrak{a}^\circ = \mathfrak{a}$ holds for vanishing ideals of subsets of Ω . Thus, by the Stone–Weierstrass approximation theorem for locally compact spaces (cf. [Con90, Chapter 5, Corollary 8.3]), we can choose, for every $\varepsilon > 0$, a Laurent polynomial $h \in \mathfrak{a} \subseteq L$ that satisfies $\sup_X |f - h| \leq \frac{\varepsilon}{2}$. We can pick a multi-index $\beta \in \mathbb{N}^n$ such that $p := x^\beta h \in R$ is a polynomial. As $|f(\xi)| = 1$, this implies in particular that

$$\frac{|p(\xi)|}{\|p\|_\infty} = \frac{|h(\xi)|}{\|h\|_\infty} \geq \frac{1 - \frac{\varepsilon}{2}}{1 + \frac{\varepsilon}{2}} \geq 1 - \varepsilon. \quad (3.12)$$

For any $d \in \mathbb{N}$, denote by $N_d := (d+1)^n$ the dimension of W_d . Recall that, for $e_{d\xi} = \sum_{|\alpha|_\infty \leq d} \frac{\xi^{-\alpha}}{\sqrt{N_d}} \frac{x^\alpha}{\sqrt{N_d}} \in W_d$, we have $N_d \int_\Omega |e_{d\xi}(x)|^2 dx = 1$ by (3.11).

Fix $\delta \in \mathbb{N}$ such that $p \in \mathfrak{a} \cap R_{\leq \delta}$. For every $d \in \mathbb{N}$, we define the polynomial

$$q_d := \frac{\sqrt{N_d}}{\sqrt{N_{d+\delta}}} e_{d\xi} p \in \mathfrak{a} \cap W_{d+\delta}.$$

Then it follows from (3.11) that

$$\|q_d\|_2^2 = N_{d+\delta} \int_\Omega |q_d(x)|^2 dx \leq N_d \int_\Omega |e_{d\xi}(x)|^2 dx \|p\|_\infty^2 = \|p\|_\infty^2.$$

On the other hand, as $e_{d\xi}(\xi) = 1$, we have

$$|q_d(\xi)|^2 = \frac{N_d}{N_{d+\delta}} |p(\xi)|^2.$$

Hence, with (3.12), it follows that

$$\liminf_{d \rightarrow \infty} 1 - P_{d+\delta,1}(\xi) \geq \liminf_{d \rightarrow \infty} \frac{|q_d(\xi)|^2}{\|q_d\|_2^2} \geq \liminf_{d \rightarrow \infty} \frac{N_d |p(\xi)|^2}{N_{d+\delta} \|p\|_\infty^2} \geq (1 - \epsilon)^2.$$

As we can choose ϵ arbitrarily small and since $1 - P_{d,1}(\xi) \leq 1$ for all $d \in \mathbb{N}$ by Lemma 3.5.7 (1), it follows that $\lim_{d \rightarrow \infty} 1 - P_{d,1}(\xi) = 1$.

The statement about $Q_{d,0}$ follows immediately from (3.10). By Lemma 3.5.10 (1), this implies that $\limsup_{d \rightarrow \infty} Q_{d,\epsilon}(\xi) \leq \epsilon$. \square

Remark 3.5.14. Assume that μ is a measure on \mathbf{T}^1 such that $\log \mu' \in L^1(\mathbf{T}^1)$, where μ' is defined as the derivative of the function $t \mapsto \int_{[0,t)} d\mu(\vartheta)$, for $t \in \mathbf{T}^1$. For instance, this holds if μ has a continuous density that is strictly positive everywhere. With the variational definition (3.8) of the Christoffel function Q_d , it then holds by [MNT91, Theorem 1] that $\lim_{d \rightarrow \infty} NQ_d(\iota(t)) = \mu'(t)$ for Lebesgue-almost all $t \in \mathbf{T}^1$. Thus, $NQ_d \circ \iota$ approximates the measure μ , for large $d \in \mathbb{N}$. Under slightly stronger assumptions, a similar statement for compactly-supported measures in higher dimensions is given in [PPL21, Theorem 4], which is stated in relation to a reference measure.

By Proposition 3.5.11, the function $Q_{d,\epsilon}$ interpolates the weights of any finitely-supported measure. Additionally, Theorem 3.5.13 shows that, everywhere outside the Zariski closure of the support, the function $Q_{d,\epsilon}$ becomes small for large $d \in \mathbb{N}$ and small $\epsilon > 0$.

Recall that $Q_{d,\epsilon} = Q_d$ holds for measures with Zariski-dense support in \mathbf{T}^n (cf. Remark 3.5.5). These observations motivate studying the function $Q_{d,\epsilon}$ (with a suitable normalization) as an approximation to a measure, even in case of measures that are supported on a positive-dimensional variety. See also [PPL21] for related results. \diamond

Example 3.5.15. We consider the uniform measure on the union of three trigonometric algebraic curves in \mathbf{T}^2 , each of which is generated by a polynomial of max-degree 1. Figure 3.1 displays the corresponding functions $P_{d,1}$, P_d and $Q_{d,\epsilon}$ for different degrees $d \in \mathbb{N}$, normalized such that their supremums are equal. For $d \rightarrow \infty$, the function $P_{d,1}$ converges to the indicator function of the whole variety on the torus, by Theorem 3.5.13. Note that $P_{d,1}$ is only non-trivial if $d \geq 3$, as the variety is generated by a single polynomial of max-degree 3. We observe that, for small $d \in \mathbb{N}$ and $\epsilon > 0$, the function $Q_{d,\epsilon}$ seems to be better localized at the variety than $P_{d,1}$ and P_d . However, our computations suggest that, in contrast to $P_{d,1}$, the normalized function $\frac{Q_{d,\epsilon}}{\|Q_{d,\epsilon}\|_\infty}$, for $d \rightarrow \infty$, does not in general converge to a function that is constant when restricted to the underlying variety, not even when the moments come from the uniform measure supported on a variety that is smooth on the torus. Nevertheless, $Q_{d,\epsilon}$ still provides a useful approximate depiction of the variety. This is in line with the convergence statement of level sets proved in [LP19, Theorem 3.9] for the Tikhonov regularization of the Christoffel function. \diamond

Example 3.5.16. On the one-dimensional torus \mathbf{T}^1 , we consider the finitely-supported measure $\sum_{j=1}^3 \lambda_j \delta_{\vartheta_j}$ defined by

$$(\vartheta_1, \lambda_1) = (0.2, 1.3), \quad (\vartheta_2, \lambda_2) = (0.35, 0.7), \quad (\vartheta_3, \lambda_3) = (0.8, 1).$$

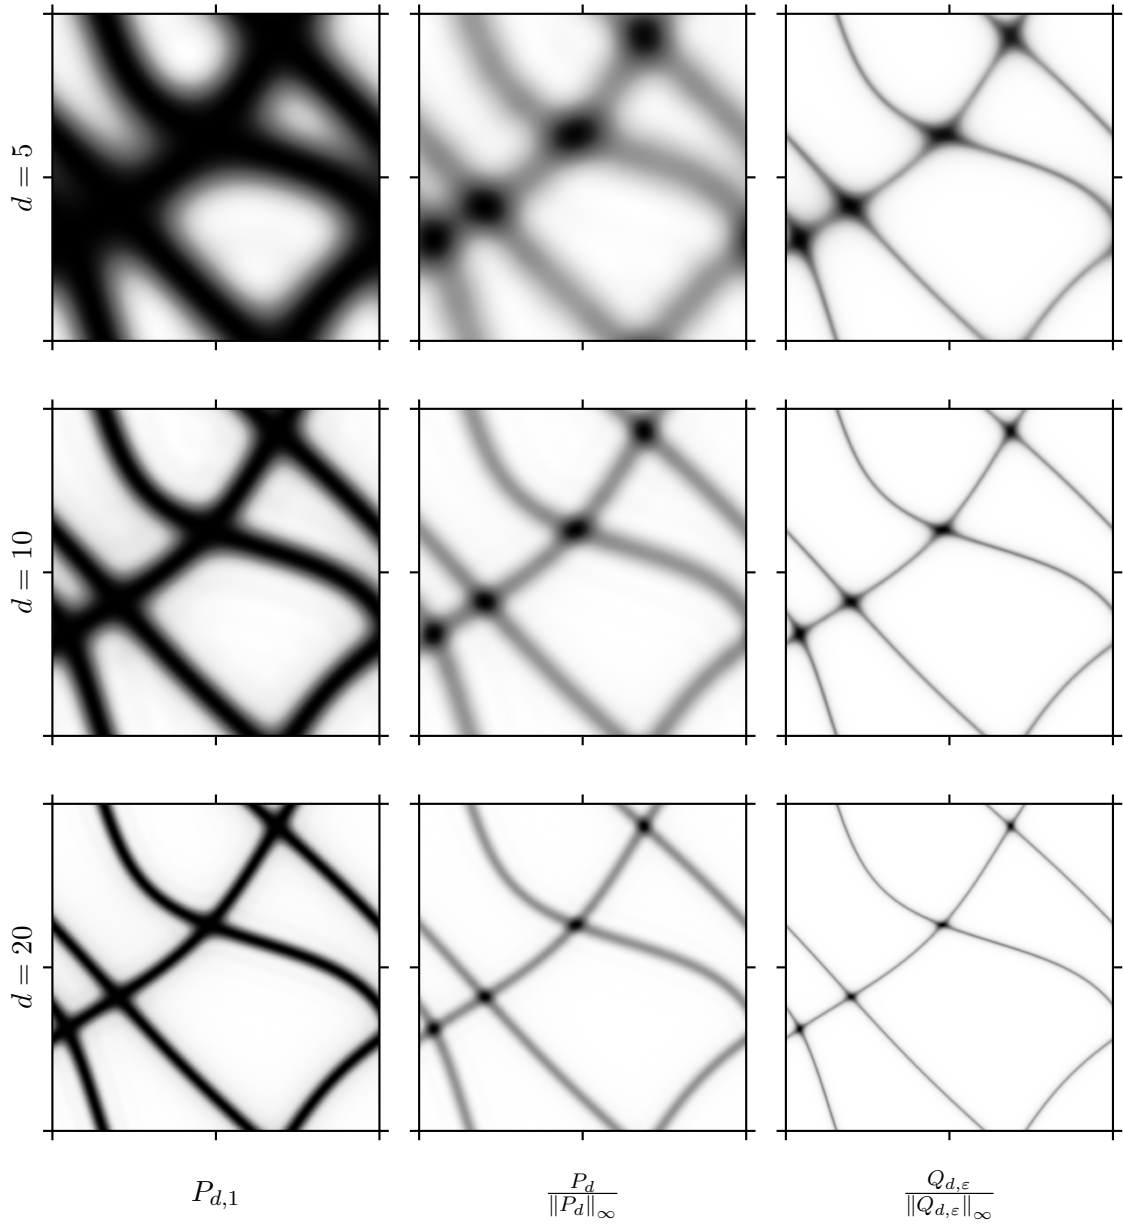


Figure 3.1: The functions $P_{d,1}$, P_d and $Q_{d,\epsilon}$, $\epsilon = 0.01$, associated to the uniform measure on the union of three trigonometric curves in \mathbf{T}^2 , for different degrees $d \in \mathbb{N}$.




Figure 3.2: The functions $P_{d,1}$ (dotted), P_d (dashed) and $Q_{d,\varepsilon_1}, Q_{d,\varepsilon_2}$ for $\varepsilon_1 := 0.1$, $\varepsilon_2 := 0.01$ (solid, dash-dotted) associated to the measure $\sum_{j=1}^3 \lambda_j \delta_{\vartheta_j}$ on \mathbf{T}^1 . The Diracs δ_{ϑ_j} are marked by \circ , the weighted Diracs $\lambda_j \delta_{\vartheta_j}$ by \bullet .

The functions $P_{d,1}$, P_d and $Q_{d,\varepsilon}$, for two different choices of $\varepsilon > 0$, are displayed in Figure 3.2. The image agrees with our previous observations. The function $P_{d,1}$ peaks at the points ϑ_j , attaining the value 1, as stated in Lemma 3.5.7. The function $Q_{d,\varepsilon}$, in contrast, does not have local maxima exactly at the points ϑ_j , but instead interpolates the weights λ_j at ϑ_j , for $1 \leq j \leq 3$, by Proposition 3.5.11. As the degree d increases, the location of the support points ϑ_j of the measure becomes clearer. Indeed, away from the points, the function $Q_{d,\varepsilon}$ tends to ε , while $P_{d,1}$ goes to 0, which agrees with Theorem 3.5.13. The behavior of P_d is explained by Lemma 3.5.19 below. \diamond

Definition 3.5.17. The univariate Fejér kernel F_d of degree $d \in \mathbb{N}$ (cf. [Gra14, Definition 3.1.8]) is defined as

$$F_d: \mathbf{T} \longrightarrow \mathbb{R}_{\geq 0}, \quad t \longmapsto \begin{cases} d+1 & \text{if } t = 0, \\ \frac{\sin^2((d+1)\pi t)}{(d+1)\sin^2(\pi t)} & \text{otherwise,} \end{cases}$$

which is continuous on \mathbf{T} . For $n \geq 2$, the multivariate Fejér kernel is defined by ten-

sorization as $F_d(t) := \prod_{i=1}^n F_d(t_i)$ for $t \in \mathbf{T}^n$. Note that, with $N := (d+1)^n$, we have

$$\begin{aligned} F_d(t) &= \prod_{i=1}^n \sum_{\alpha_i=-d}^d \left(1 - \frac{|\alpha_i|}{d+1}\right) e^{2\pi i \alpha_i t_i} \\ &= \frac{1}{N} \prod_{i=1}^n \left| \sum_{\alpha_i=0}^d e^{2\pi i \alpha_i t_i} \right|^2 = \frac{1}{N} \left| \sum_{\alpha \in \mathbb{N}^n, |\alpha|_\infty \leq d} e^{2\pi i \langle \alpha, t \rangle} \right|^2 \end{aligned} \quad (3.13)$$

for all $t \in \mathbf{T}^n$.

Definition 3.5.18. For a function f and a (signed or non-negative) measure μ on \mathbf{T}^n , the *convolution* $f * \mu$ is a function given by

$$(f * \mu)(\vartheta) = \int_{\mathbf{T}^n} f(\vartheta - t) d\mu(t),$$

for $\vartheta \in \mathbf{T}^n$ (cf. [Sch73, Chapter 6]). It satisfies

$$\int_{\mathbf{T}^n} \varphi d(f * \mu) = \int_{\mathbf{T}^n} (\check{f} * \varphi) d\mu,$$

for all continuous functions φ on \mathbf{T}^n , where \check{f} denotes the *reflection* of f , i. e. $\check{f}(t) := f(-t)$ for all $t \in \mathbf{T}^n$. Recall that $(\check{f} * \varphi)(\vartheta) = \int_{\mathbf{T}^n} \check{f}(\vartheta - t) \varphi(t) dt$, for all $\vartheta \in \mathbf{T}^n$.

Lemma 3.5.19. *Let μ be a measure on \mathbf{T}^n . Then $NP_d \circ \iota = F_d * \mu$, for $d \in \mathbb{N}$ and $N := (d+1)^n$. In particular*

$$\lim_{d \rightarrow \infty} N \int_{\mathbf{T}^n} \varphi(t) P_d(\iota(t)) dt = \int_{\mathbf{T}^n} \varphi(t) d\mu(t),$$

for any continuous function φ on \mathbf{T}^n .

This property is also called *weak* convergence* of $NP_d \circ \iota$ with limit μ for $d \rightarrow \infty$ (cf. [Kad18, Section 17.2]), which we denote by

$$NP_d \circ \iota \xrightarrow[d \rightarrow \infty]{\text{weak}^*} \mu.$$

Proof. By the choice of B_d and by (3.13), we have

$$|e_{d\xi}(x)|^2 = \frac{1}{N^2} \left| \sum_{\alpha \in \mathbb{N}^n, |\alpha|_\infty \leq d} (x\bar{\xi})^\alpha \right|^2 = \frac{1}{N} F_d(t - \vartheta), \quad (3.14)$$

where $x = \iota(t)$ and $\xi = \iota(\vartheta)$, $\vartheta \in \mathbf{T}^n$. We therefore obtain that

$$NP_d(\iota(\vartheta)) = N e_{d\xi}^* H_d e_{d\xi} = \int_{\mathbf{T}^n} F_d(t - \vartheta) d\mu(t),$$

for all $\vartheta \in \mathbf{T}^n$, which shows that $NP_d \circ \iota = F_d * \mu$, since $\tilde{F}_d = F_d$.

The addendum then holds due to the weak* convergence of $F_d * \mu$ to μ for $d \rightarrow \infty$. Indeed, as $\tilde{F}_d = F_d$, for a continuous function φ on \mathbf{T}^n it holds that

$$\begin{aligned} \left| \int_{\mathbf{T}^n} \varphi d(F_d * \mu) - \int_{\mathbf{T}^n} \varphi d\mu \right| &= \left| \int_{\mathbf{T}^n} (F_d * \varphi) d\mu - \int_{\mathbf{T}^n} \varphi d\mu \right| \\ &\leq \int_{\mathbf{T}^n} |(F_d * \varphi)(t) - \varphi(t)| d\mu(t) \\ &\leq \|F_d * \varphi - \varphi\|_\infty \int_{\mathbf{T}^n} d\mu(t) \end{aligned}$$

and $\lim_{d \rightarrow \infty} \|F_d * \varphi - \varphi\|_\infty = 0$ by [Gra14, Theorem 1.2.19, Proposition 3.1.10]. \square

Theorem 3.5.20. *Let μ be a finitely-supported measure on \mathbf{T}^n and let $N := (d + 1)^n$ for $d \in \mathbb{N}$. Then*

$$NP_{d,1} \circ \iota \xrightarrow[d \rightarrow \infty]{\text{weak}^*} \sum_{\vartheta \in \text{supp } \mu} \delta_\vartheta,$$

so, for any continuous function φ on \mathbf{T}^n , we have

$$\lim_{d \rightarrow \infty} N \int_{\mathbf{T}^n} \varphi(t) P_{d,1}(\iota(t)) dt = \sum_{\vartheta \in \text{supp } \mu} \varphi(\vartheta).$$

Proof. For sufficiently large $d \in \mathbb{N}$, the rank r of the moment matrix H_d is equal to the number of support points of μ . We denote the distinct support points by $\vartheta_1, \dots, \vartheta_r \in \mathbf{T}^n$ and set $\xi_j := \iota(\vartheta_j)$ for $1 \leq j \leq r$. Next, we define the function $a(t) := \sum_{j=1}^r F_d(t - \vartheta_j)$ for $t \in \mathbf{T}^n$. Then

$$\begin{aligned} &\left| N \int_{\mathbf{T}^n} \varphi(t) P_{d,1}(\iota(t)) dt - \sum_{j=1}^r \varphi(\vartheta_j) \right| \\ &\leq \left| \int_{\mathbf{T}^n} \varphi(t) (NP_{d,1}(\iota(t)) - a(t)) dt \right| + \left| \int_{\mathbf{T}^n} \varphi(t) a(t) dt - \sum_{j=1}^r \varphi(\vartheta_j) \right| \\ &\leq \|NP_{d,1} \circ \iota - a\|_1 \cdot \|\varphi\|_\infty + \left| \int_{\mathbf{T}^n} \varphi(t) a(t) dt - \sum_{j=1}^r \varphi(\vartheta_j) \right|, \end{aligned}$$

where the latter inequality follows from Hölder's inequality. Since $a = F_d * \sum_{j=1}^r \delta_{\vartheta_j}$, the term on the right converges to 0 for $d \rightarrow \infty$ by Lemma 3.5.19. Therefore, it is enough to show that $\lim_{d \rightarrow \infty} \|NP_{d,1} \circ \iota - a\|_1 = 0$.

For this, we define the $r \times N$ -matrix

$$A = (w(\xi_j))_{1 \leq j \leq r, w \in B_d}$$

which, up to a scalar factor, is a Vandermonde matrix. With Equation (3.14), it follows that

$$a(t) = \sum_{j=1}^r F_d(t - \vartheta_j) = \sum_{j=1}^r N |e_{d\xi_j}(x)|^2 = N e_{dx}^* A^* A e_{dx},$$

where $x = \iota(t)$. If d is sufficiently large such that H_d has rank r , note further that

$$U_1 \operatorname{diag}(\varsigma_1, \dots, \varsigma_r) U_1^* = H_d = \left(\sum_{j=1}^r \lambda_j \overline{v(\xi_j)} w(\xi_j) \right)_{v, w \in B_d} = A^* \operatorname{diag}(\lambda_1, \dots, \lambda_r) A,$$

for some weights $\lambda_1, \dots, \lambda_r \in \mathbb{C}^*$. As then A must be a matrix of rank r , the images of U_1 and A^* are the same, so we can assume that $A^* = U_1 C$ for some invertible $r \times r$ -matrix C . Thus, we obtain

$$\begin{aligned} |a(t) - N P_{d,1}(\iota(t))| &= N |e_{dx}^* (A^* A - U_1 U_1^*) e_{dx}| = N |e_{dx}^* U_1 (C C^* - I_r) U_1^* e_{dx}| \\ &\leq N \|U_1^* e_{dx}\|_2^2 \cdot \|C C^* - I_r\|_2 \quad (3.15) \\ &= N P_{d,1}(\iota(t)) \cdot \|C C^* - I_r\|_2, \end{aligned}$$

where $x = \iota(t)$.

Since the eigenvalues of the matrices $C C^*$ and $C^* C$ are the same by [HJ13, Theorem 1.3.22], also the matrices $C C^* - I_r$, $C^* C - I_r$ have equal eigenvalues. As these matrices are Hermitian and thus normal, their singular values are equal to the absolute values of their eigenvalues (cf. [HJ13, Problem 2.6.P15]), so we deduce that

$$\|C C^* - I_r\|_2 = \|C^* C - I_r\|_2. \quad (3.16)$$

Further, observe that $AA^* = (\sum_{w \in B_d} w(\xi_j) \overline{w(\xi_l)})_{1 \leq j, l \leq r}$. In particular, the diagonal entries of this matrix are equal to 1, by the addition theorem (3.6). Additionally, denote the minimal separation distance between the points on the torus by $\vartheta_{\min} := \min_{1 \leq j \neq l \leq r} \max_{1 \leq i \leq n} |\vartheta_{ji} - \vartheta_{li}|$, where $|\cdot|$ denotes the induced metric on \mathbf{T}^1 . Note that $0 < \vartheta_{\min} \leq \frac{1}{2}$, so $\sin(\pi \vartheta_{\min}) \geq 2\vartheta_{\min}$ and therefore

$$\frac{1}{N} F_d(\vartheta_j - \vartheta_l) \leq \frac{(d+1)^{n-1}}{N} \cdot \frac{1}{(d+1) \sin^2(\pi \vartheta_{\min})} \leq \frac{1}{(d+1)^2 (2\vartheta_{\min})^2},$$

for $j \neq l$. Together with (3.16), we then obtain the upper bound

$$\begin{aligned} \|C C^* - I_r\|_2 &= \|C^* C - I_r\|_2 = \|A A^* - I_r\|_2 \\ &\leq \|A A^* - I_r\|_F = \left(\sum_{1 \leq j \neq l \leq r} \left| \sum_{w \in B_d} w(\xi_j) \overline{w(\xi_l)} \right|^2 \right)^{\frac{1}{2}} = \left(\sum_{1 \leq j \neq l \leq r} \frac{1}{N} F_d(\vartheta_j - \vartheta_l) \right)^{\frac{1}{2}} \\ &\leq \frac{\sqrt{r(r-1)}}{(d+1)2\vartheta_{\min}}, \end{aligned}$$

which converges to 0 for $d \rightarrow \infty$. By (3.15), this implies that

$$\|NP_{d,1} \circ \iota - a\|_1 \leq \frac{\sqrt{r(r-1)}}{(d+1)2\vartheta_{\min}} \|NP_{d,1} \circ \iota\|_1.$$

Since $\|NP_{d,1} \circ \iota\|_1 = r$ is constant by Equation (3.11), the result follows. \square

Remark 3.5.21. The proof of Theorem 3.5.20 does not use the fact that the weights $\lambda_1, \dots, \lambda_r$ are positive. Hence, the statement can be extended to signed complex measures, as explained in Remark 3.5.8. \diamond

4 Recovery of components from eigenvalues

In this chapter, we develop tools for recovering the underlying components of a mixture of measures supported on algebraic varieties. For zero-dimensional algebraic varieties, i. e. for finitely-supported measures, this problem is often addressed by solving generalized eigenproblems. Here, we study to which extent such an eigenvalue-based approach can be transferred to the case of measures supported on algebraic varieties of any dimension.

After a brief exposition and motivation of the topic in Section 4.1, we start with a thorough introduction to (generalized) eigenvalues and eigenvectors of matrix pencils consisting of a family of matrices $(\Delta_0, \dots, \Delta_r)$ in Section 4.2, which is central for our analysis. We collect several results about matrix pencils that are relevant for the rest of the chapter.


In Section 4.3, we then give a first recovery algorithm, Algorithm 4.1, which is stated generally for linear transformations of matrix pencils, i. e. without referring to moment problems. This algorithm is not guaranteed to manage to reconstruct the components in all cases, but it serves expository purposes of the eigenvalue-based recovery approach. We illustrate the algorithm and its shortcomings with moment problems on zero-dimensional and positive-dimensional varieties. The algorithms in later sections build upon it, guaranteeing successful reconstruction under some adaptations.

One such algorithm is presented in Section 4.4. In this section, we focus on positive-semidefinite matrices and eigenvalues with non-negative coordinates, in which case additional conclusions are possible due to convexity of the problem.

We then continue in Section 4.5 with an emphasis on moment problems on algebraic varieties. By making full use of the algebraic structure of the problem we are interested in, we develop a criterion that allows us to overcome the drawbacks of Algorithm 4.1. As a brief interlude in Section 4.6, we illustrate by examples that the criterion developed in the previous section is sharp. In Section 4.7, the criterion is then employed in two improved algorithms. These are variants of Algorithm 4.1, but stated specifically for moment problems on varieties, in which case full recovery is possible.

We finish with a discussion of details that are important for an implementation of the algorithms and with several concrete numerical examples, in Section 4.8.

The following diagram displays the suggested reading order of the sections, some of which can be considered optional on a first reading of the chapter.



4.1 Prologue

In case of a finitely-supported (non-negative or signed) measure, it is possible to compute the support points from the generalized eigenvalues of matrix pencils of shifted moment matrices, as we outline below. This is the foundation of established pencil-based recovery methods such as ESPRIT. See for instance [HS90; RK90; ACdH10; Moi15]. Such a measure is a mixture of measures supported at single points (which in particular are *zero-dimensional* algebraic varieties, to which we refer as the *components* of the measure).

In this chapter, we want to consider measures that are mixtures of measures supported on algebraic varieties that are allowed to be *positive-dimensional*. In the prototypical example we have in mind, the individual components are varieties that are pairwise not contained in each other, but differing dimensions are permitted. In this setting, the eigenvalues of matrix pencils of shifted moment matrices are not enough for recovering the underlying components.

Assume that $\sigma_j: \mathbb{k}[x_1, \dots, x_n] \rightarrow \mathbb{k}$ are the moment functionals of different components and $\sigma = \sum_{j=1}^r \lambda_j \sigma_j$ is the moment functional of the mixture, where $\lambda_1, \dots, \lambda_r \in \mathbb{k}^*$. For some $d \in \mathbb{N}$, define the moment matrix $H = (\sigma(x^{\alpha+\beta}))_{|\alpha|, |\beta| \leq d}$ and the shifted matrices $H^{(i)} = (\sigma(x_i x^{\alpha+\beta}))_{|\alpha|, |\beta| \leq d}$ for $1 \leq i \leq n$. We consider the matrix pencil

$$H^{(i)} - \gamma H = \sum_{j=1}^r \lambda_j (\sigma_j((x_i - \gamma)x^{\alpha+\beta}))_{|\alpha|, |\beta| \leq d},$$

for $1 \leq i \leq n$. Its finite eigenvalues are the values $\gamma \in \mathbb{k}$ that are rank-reducing, as explained in more detail in Section 4.2.

In the zero-dimensional case, we can assume that the functionals are of the form $\sigma_j = \text{ev}_{\xi_j}$ for distinct points $\xi_j \in \mathbb{k}^n$, $1 \leq j \leq r$. Then the above pencil is of the form

$$H^{(i)} - \gamma H = \sum_{j=1}^r \lambda_j (\xi_{ji} - \gamma) (\xi_j^{\alpha+\beta})_{|\alpha|, |\beta| \leq d},$$

where the matrices $(\xi_j^{\alpha+\beta})_{|\alpha|,|\beta|\leq d}$ are of rank 1. If d is sufficiently large, γ is rank-reducing if and only if $\gamma = \xi_{ji}$ for some $j \in \{1, \dots, r\}$, so the eigenvalues of the pencil $H^{(i)} - \gamma H$ are exactly the i -th coordinates of the points ξ_1, \dots, ξ_r , which thus presents a direct method for recovering the coordinates of the support points. We will make this more explicit in Example 4.3.4. In particular, we will assign meaning to the corresponding eigenvectors, as well.

While this works in case of zero-dimensional components, this is not usually possible for positive-dimensional components, so there is no γ such that any of the matrices $(\sigma_j((x_i - \gamma)x^{\alpha+\beta}))_{|\alpha|,|\beta|\leq d}$, $1 \leq j \leq r$, has smaller rank than for generic choices of $\gamma \in \mathbb{k}$. An intuitive reason for this is the well-known fact that, in the zero-dimensional case, the shifted moment matrices give rise to commuting multiplication operators (see e.g. [Mou18]) that are simultaneously diagonalizable (cf. [Mey00, Exercise 7.2.16] and Remark 4.2.10) – this fails to hold in the positive-dimensional case, since the coordinate ring of the underlying positive-dimensional variety is not a finite-dimensional vector space.

Instead of considering pencils involving the shifted moment matrices $H^{(i)}$, $1 \leq i \leq n$, in this chapter we consider a family (M_0, \dots, M_s) of moment matrices which are of the form $M_k = \sum_{j=1}^r \lambda_{kj} H_j$, $0 \leq k \leq s$, for suitable coefficients $\lambda_{kj} \in \mathbb{k}$, where $H_j = (\sigma_j(x^{\alpha+\beta}))_{|\alpha|,|\beta|\leq d}$ are the moment matrices of individual components. We assume that the matrix $(\lambda_{kj})_{kj} \in \mathbb{k}^{(s+1) \times r}$ has full column rank. With input of this kind, we will see that an eigenvalue-based approach can be used to recover the moment matrices H_j , $1 \leq j \leq r$, of the individual components. From there, one can use Theorem 3.4.11 to recover the individual vanishing ideals of these components. In Proposition 4.3.11 and Remark 4.3.13 as well as Example 4.3.16, we relate this new approach to the conventional method that is based on the shifted moment matrices, as outlined above. A fundamental reason for why this can work is the observation that there exist eigenvectors which can be interpreted as polynomials that vanish on all but one of the components.

Data of the form described above arises, for example, in the field of multi-snapshot spectral estimation (see e.g. [LZGL21]), such as direction-of-arrival estimation and time series analysis [KV96]. In this context, one considers a time-dependent measure

$$\nu(t) = \sum_{j=1}^r L_j(t) \mu_j, \quad t \in \mathbb{R},$$

for some functions $L_j: \mathbb{R} \rightarrow \mathbb{k}$, $1 \leq j \leq r$, where μ_j is some measure (such as a Dirac measure) with moment functional σ_j for $1 \leq j \leq r$. The matrices (M_0, \dots, M_s) above can then be interpreted as the moment matrices of the measure $\nu(t)$ at multiple different timestamps $t_0, \dots, t_s \in \mathbb{R}$ by setting $\lambda_{kj} := L_j(t_k)$ for all $0 \leq k \leq s$. (The stability results of [LZGL21] for the univariate case of finitely-supported measures suggest that our assumption that the matrix $(\lambda_{kj})_{kj} \in \mathbb{k}^{(s+1) \times r}$ is of rank r is a natural one.) Our analysis extends the classic scenario, in which the measures μ_j are Diracs, to more general measures supported on varieties of any dimension, in a noise-free setting.

4.2 Generalized eigenvalues

Here, we introduce the notion of generalized eigenvalues of a regular pencil of multiple matrices $(\Delta_0, \dots, \Delta_r)$, which is central for the rest of the chapter. The main reference on this topic is [Atk72]. The statements in this section hold for an arbitrary field \mathbb{k} .

4.2.1 Definitions and elementary properties

Definition 4.2.1 ([Atk72, Chapter 6], [MP09], [Car21]). Let $\Delta_0, \dots, \Delta_r \in \mathbb{k}^{n \times n}$, $r \geq 1$, be square matrices. Then the matrix pencil $(\Delta_0, \dots, \Delta_r)$ is called *regular* if there exists a linear combination

$$\Delta = \sum_{j=0}^r \lambda_j \Delta_j \quad (4.1)$$

with $\lambda_0, \dots, \lambda_r \in \mathbb{k}$ such that Δ is an invertible matrix. Otherwise, the pencil is called *singular*. Thus, if \mathbb{k} is an infinite field, a pencil is regular if and only if

$$\det \left(\sum_{j=0}^r z_j \Delta_j \right) \neq 0 \quad (4.2)$$

as a polynomial in the variables z_0, \dots, z_r . (This does not in general hold over a finite field, as there exist non-zero polynomials that vanish everywhere.)

For a regular matrix pencil $(\Delta_0, \dots, \Delta_r)$, a point $\gamma = [\gamma_0 : \dots : \gamma_r] \in \mathbb{P}_{\mathbb{k}}^r$ is a (*generalized*) *eigenvalue* of the pencil if there exists a vector $v \neq 0$ in

$$\bigcap_{j=0}^r \ker(\Delta^{-1} \Delta_j - \gamma_j \mathbf{I}_n), \quad (4.3)$$

where Δ is an invertible matrix as in (4.1). In this case, v is a (*generalized*) *eigenvector* of the pencil with corresponding eigenvalue γ . Thus, the eigenvector and eigenvalue satisfy

$$\Delta_j v = \gamma_j \Delta v$$

for all $0 \leq j \leq r$.

The associated eigenvalue problem is also referred to as *coupled* eigenvalue problem. As there is no risk of confusion, we often refer to the generalized eigenvalues and generalized eigenvectors simply as eigenvalues and eigenvectors of the matrix pencil. Though, it is important to note that a generalized eigenvector of the pencil $(\Delta_0, \dots, \Delta_r)$ is not an ordinary eigenvector of each of the matrices $\Delta_0, \dots, \Delta_r$, but it is a simultaneous ordinary eigenvector of all the matrices $\Delta^{-1} \Delta_0, \dots, \Delta^{-1} \Delta_r$. To this end, see also Remark 4.2.10.

Remark 4.2.2. For pencils (A, B) of two matrices $A, B \in \mathbb{k}^{n \times n}$, the theory of generalized eigenvalues is much more developed, including the singular case. See [Dem00] for an overview. It is very common to write such a pencil as $A - zB$, where z denotes a variable. For the eigenvalues, one then chooses inhomogeneous coordinates. So, assuming the pencil is regular, $\gamma \in \mathbb{k} \cup \{\infty\}$ is an eigenvalue if γ is a root of the polynomial $\det(A - zB)$ in case γ is finite or if the polynomial has degree smaller than n in case γ is infinite. In terms of homogeneous coordinates, the eigenvalue $\gamma \in \mathbb{k} \cup \{\infty\}$ of the pencil (A, B) is $[\gamma : 1] \in \mathbb{P}_{\mathbb{k}}^1$ if $\gamma \neq \infty$ and $[1 : 0] \in \mathbb{P}_{\mathbb{k}}^1$ if $\gamma = \infty$.

It is possible to extend the definition of eigenvalues to singular matrix pencils (see, e.g. [Sle+00, Chapter 8.7] and [MP09]), but here we focus on regular matrix pencils. It is worth noting that the notion of eigenvector, however, is not well-defined for singular matrix pencils. Instead, one can define *reducing subspaces*, which is a generalization of the concept of eigenspaces as well as invariant or deflating subspaces; see [Van83]. \diamond

Remark 4.2.3. Since every coordinate γ_j is an ordinary eigenvalue of the matrix $\Delta^{-1}\Delta_j$ for $0 \leq j \leq r$, it is clear that a regular pencil $(\Delta_0, \dots, \Delta_r)$ can only have finitely many distinct eigenvalues $\gamma = [\gamma_0 : \dots : \gamma_r]$ if the eigenvalue corresponding to a given eigenvector is well-defined. The latter is proved in Lemma 4.2.8. A stronger bound on the number of eigenvalues is given in Lemma 4.2.9.

Furthermore, note that an eigenvalue γ is only determined up to scaling. The representatives $\gamma_0, \dots, \gamma_r$ depend on the choice of $\lambda_0, \dots, \lambda_r$. Indeed, if $(\gamma_0, \dots, \gamma_r)$ satisfies (4.3) for our choice of Δ , then the condition is also satisfied by $c(\gamma_0, \dots, \gamma_r)$ in terms of the invertible matrix $c^{-1}\Delta$, for any non-zero scalar $c \in \mathbb{k}$. This justifies defining the eigenvalue as the projective class $[\gamma_0 : \dots : \gamma_r] \in \mathbb{P}_{\mathbb{k}}^r$. The fact that the eigenvectors do not depend on the choice of Δ follows from Proposition 4.2.5 below.

Additionally, note that we have

$$\gamma_i \gamma_j \Delta v = \gamma_i \Delta_j v = \gamma_j \Delta_i v$$

for $i \neq j$, if $[\gamma_0 : \dots : \gamma_r]$ is an eigenvalue with eigenvector v . If (Δ_i, Δ_j) is a regular pencil, the latter equation implies that v is an eigenvector of the pencil (Δ_i, Δ_j) with eigenvalue $[\gamma_i : \gamma_j]$ if $\gamma_i, \gamma_j \neq 0$. However, this need not hold in general, and indeed, in many cases of interest, (Δ_i, Δ_j) is a singular matrix pencil for some or all pairs $i \neq j$, even though the pencil $(\Delta_0, \dots, \Delta_r)$ is regular. For instance, this is the case in Example 4.2.4 below.

A method for computing the generalized eigenvalues and eigenvectors of a regular pencil, based on the computation of a QZ-decomposition, also called generalized Schur form, is described in [HKP04]. A variant thereof is described in more detail in Remark 4.8.9, where we discuss a concrete implementation. \diamond

Example 4.2.4. Consider the matrix pencil

$$(\Delta_0, \Delta_1, \Delta_2) = (\text{diag}(1, 0, 0), \text{diag}(0, 1, 0), \text{diag}(0, 0, 1)).$$

Then $\Delta_0 + \Delta_1 + \Delta_2$ is the identity matrix, so the pencil is regular, but each of the pencils (Δ_0, Δ_1) , (Δ_0, Δ_2) and (Δ_1, Δ_2) is singular. \diamond

Proposition 4.2.5. *Let $(\Delta_0, \dots, \Delta_r)$ be a regular matrix pencil and let v be a non-zero vector. Then the following are equivalent:*

- (1) *The vector v is an eigenvector of the pencil $(\Delta_0, \dots, \Delta_r)$.*
- (2) *The subspace*

$$\langle \Delta_0 v, \dots, \Delta_r v \rangle$$

is one-dimensional.

- (3) *There exists a point $\gamma = [\gamma_0 : \dots : \gamma_r] \in \mathbb{P}_{\mathbb{k}}^r$ such that*

$$\gamma_i \Delta_j v = \gamma_j \Delta_i v$$

holds for all $0 \leq i < j \leq r$. In this case, γ is the corresponding eigenvalue.

Proof. As the pencil is regular, we may assume that Δ is a linear combination of $\Delta_0, \dots, \Delta_r$ which is an invertible matrix. If $v \neq 0$ is an eigenvector of the pencil with eigenvalue $[\gamma_0 : \dots : \gamma_r]$, we may assume that the representatives γ_j are chosen such that

$$\Delta_j v = \gamma_j \Delta v$$

holds for all $0 \leq j \leq r$. Hence, we have $\langle \Delta_0 v, \dots, \Delta_r v \rangle = \langle \Delta v \rangle$, as not all γ_j are zero. Since Δ is invertible and $v \neq 0$, this space is one-dimensional, showing that (1) implies (2).

In order to prove that (2) implies (3), we assume that $\langle \Delta_0 v, \dots, \Delta_r v \rangle = \langle w \rangle$ is spanned by some non-zero vector w . We choose γ_j such that $\Delta_j v = \gamma_j w$ holds for $0 \leq j \leq r$, so it follows that

$$\gamma_i \Delta_j v = \gamma_i \gamma_j w = \gamma_j \Delta_i v$$

for all i, j . As $\langle \gamma_0 w, \dots, \gamma_r w \rangle = \langle w \rangle$, we have $\gamma_j \neq 0$ for at least one j , so $[\gamma_0 : \dots : \gamma_r] \in \mathbb{P}_{\mathbb{k}}^r$.

Assuming that $\gamma_i \Delta_j v = \gamma_j \Delta_i v$ holds for all $0 \leq i, j \leq r$ as in (3), denote by $\Delta' = \sum_{j=0}^r \lambda_j \Delta_j$ a non-singular matrix for a suitable choice of $\lambda_0, \dots, \lambda_r \in \mathbb{k}$. Then it follows that

$$\gamma_i \Delta' v = \gamma_i \sum_{j=0}^r \lambda_j \Delta_j v = \sum_{j=0}^r \lambda_j \gamma_j \Delta_i v$$

for all $0 \leq i \leq r$. As $\gamma_i \neq 0$ for some i and since Δ' is invertible and v is non-zero, this implies that $\sum_{j=0}^r \lambda_j \gamma_j \neq 0$. Consequently, we obtain

$$\frac{\gamma_i}{\sum_{j=0}^r \lambda_j \gamma_j} \Delta' v = \Delta_i v,$$

for all $0 \leq i \leq r$, so v is indeed an eigenvector with eigenvalue $[\gamma_0 : \dots : \gamma_r]$ and thus (1) holds. \square

Remark 4.2.6. The characterization of eigenvectors given in Proposition 4.2.5 (2) is sometimes used as definition of eigenvectors of a matrix pencil, even for singular pencils (see, e. g. [Rin17]), but this is non-standard terminology. We will only be concerned with regular pencils, in which case the notions agree.

The proof of Proposition 4.2.5 is similar to that of [Atk72, Theorem 6.8.1], which states that $\sum_{j=0}^r \lambda_j \gamma_j \neq 0$ holds for eigenvalues $[\gamma_0 : \dots : \gamma_r]$ of regular pencils $(\Delta_0, \dots, \Delta_r)$ where, as before, the matrix $\sum_{j=0}^r \lambda_j \Delta_j$ is non-singular. \diamond

As it is sometimes important to work with a particular choice of coordinates for the eigenvalues with respect to a fixed non-singular linear combination of the matrices of a pencil, we emphasize the following corollary, for future reference.

Corollary 4.2.7. *Let $(\Delta_0, \dots, \Delta_r)$ be a regular pencil of matrices over \mathbb{k} and let $\gamma = [\gamma_0 : \dots : \gamma_r] \in \mathbb{P}_{\mathbb{k}}^r$ be an eigenvalue of the pencil with eigenvector $v \neq 0$. If $\Delta = \sum_{j=0}^r \lambda_j \Delta_j$, with $\lambda_0, \dots, \lambda_r \in \mathbb{k}$, is any non-singular linear combination, then coordinates $(\tilde{\gamma}_0, \dots, \tilde{\gamma}_r) \in \mathbb{k}^{r+1}$ for γ are uniquely determined by the requirement $\Delta_j v = \tilde{\gamma}_j \Delta v$ and they satisfy*

$$\tilde{\gamma}_j = \frac{\gamma_j}{\sum_{k=0}^r \lambda_k \gamma_k}, \quad 0 \leq j \leq r.$$

Proof. This follows from the proof of Proposition 4.2.5. \square

Lemma 4.2.8. *Let $\Delta_0, \dots, \Delta_r \in \mathbb{k}^{n \times n}$, $r \geq 1$, be matrices such that the matrix pencil $(\Delta_0, \dots, \Delta_r)$ is regular. If $v \neq 0$ is an eigenvector of the pencil, then its eigenvalue in $\mathbb{P}_{\mathbb{k}}^r$ is unique.*

Proof. Assume that $\gamma, \delta \in \mathbb{P}_{\mathbb{k}}^r$ with $\gamma \neq \delta$ satisfy

$$\Delta_j v = \gamma_j \Delta v = \delta_j \Delta' v, \quad (4.4)$$

for all $0 \leq j \leq r$, where Δ, Δ' are suitable linear combinations of $\Delta_0, \dots, \Delta_r$ which are invertible. As $\gamma \neq \delta$, without loss of generality, we may assume that (γ_0, γ_1) and (δ_0, δ_1) are linearly independent. By Proposition 4.2.5, we have $\gamma_0 \Delta_1 v = \gamma_1 \Delta_0 v$ and $\delta_0 \Delta_1 v = \delta_1 \Delta_0 v$. This implies

$$(\delta_0 \gamma_1 - \gamma_0 \delta_1) \Delta_0 v = (\delta_0 \gamma_0 - \gamma_0 \delta_0) \Delta_1 v = 0,$$

so it follows that $\Delta_0 v = 0$, since $\delta_0 \gamma_1 - \gamma_0 \delta_1 \neq 0$ by assumption. As a consequence of (4.4), we then have $\gamma_0 = \delta_0 = 0$, as Δ and Δ' are regular matrices and $v \neq 0$. This is a contradiction to the assumption that (γ_0, γ_1) and (δ_0, δ_1) are linearly independent. \square

Lemma 4.2.9. *Let $(\Delta_0, \dots, \Delta_r)$, $r \geq 1$, be a regular pencil of matrices in $\mathbb{k}^{n \times n}$ and let $s \in \mathbb{N}$. If $\gamma_0, \dots, \gamma_s \in \mathbb{P}_{\mathbb{k}}^r$ are distinct eigenvalues of the pencil with corresponding eigenvectors $v_0, \dots, v_s \in \mathbb{k}^n$, then the eigenvectors v_0, \dots, v_s are linearly independent.*

In particular, this means that the number of eigenvalues of a pencil is bounded by the size n of the matrices.

Proof. We prove the statement by induction on s . For $s = 0$, it is correct, so let $s \geq 1$. By the inductive hypothesis, we can assume that the vectors v_1, \dots, v_s are linearly independent. Let us assume that the vectors v_0, \dots, v_s are linearly dependent, so we can write $v_0 = \sum_{i=1}^s c_i v_i$ for some $c_1, \dots, c_s \in \mathbb{k}$. Furthermore, we can choose coordinates $(\gamma_{i0}, \dots, \gamma_{ir})$ for each eigenvalue γ_i , $0 \leq i \leq s$, such that $\Delta_j v_i = \gamma_{ij} \Delta v_i$ holds for all $0 \leq j \leq r$ and some non-singular matrix Δ . From this, it follows that

$$\gamma_{0j} \sum_{i=1}^s c_i v_i = \gamma_{0j} v_0 = \Delta^{-1} \Delta_j v_0 = \Delta^{-1} \Delta_j \sum_{i=1}^s c_i v_i = \sum_{i=1}^s c_i \gamma_{ij} v_i$$

for all $0 \leq j \leq r$. As the vectors v_1, \dots, v_s are linearly independent, this implies $\gamma_{0j} c_i = c_i \gamma_{ij}$ for all $1 \leq i \leq s$ and $0 \leq j \leq r$. Without loss of generality, we can assume that $c_1 \neq 0$. Thus, it follows that $\gamma_{0j} = \gamma_{1j}$ for all $0 \leq j \leq r$, which is a contradiction to the hypothesis that the eigenvalues are distinct. \square

Remark 4.2.10. Let $(\Delta_0, \dots, \Delta_r)$, $r \geq 1$, be a regular pencil of matrices in $\mathbb{k}^{n \times n}$ and let $v_1, \dots, v_s \in \mathbb{k}^n$ be a maximal family of linearly independent eigenvectors with corresponding eigenvalues $\gamma_1, \dots, \gamma_s \in \mathbb{P}_{\mathbb{k}}^r$, which are possibly not distinct. Assume that Δ is a non-singular linear combination of the matrices such that, for all $1 \leq i \leq s$, the coordinates of γ_i satisfy $\Delta^{-1} \Delta_j v_i = \gamma_{ij} v_i$, $0 \leq j \leq r$. Denote by $Z \in \mathbb{k}^{n \times n}$ an invertible matrix whose first s columns are the eigenvectors v_1, \dots, v_s .

If $s = n$, then the eigenvectors form a basis in terms of which the regular pencil $(\Delta^{-1} \Delta_0, \dots, \Delta^{-1} \Delta_r)$ consists of diagonal matrices. Indeed, we have

$$Z^{-1} \Delta^{-1} \Delta_j Z = \text{diag}(\gamma_{1j}, \dots, \gamma_{nj})$$

for all $0 \leq j \leq r$, so the matrices $(\Delta^{-1} \Delta_0, \dots, \Delta^{-1} \Delta_r)$ are *simultaneously diagonalizable*.

Otherwise, when $s < n$, the matrices are not simultaneously diagonalizable and we obtain the block form

$$Z^{-1} \Delta^{-1} \Delta_j Z = \begin{pmatrix} \Gamma_j & * \\ 0 & * \end{pmatrix},$$

where $\Gamma_j = \text{diag}(\gamma_{1j}, \dots, \gamma_{sj})$ is a diagonal matrix, for $0 \leq j \leq r$. In this case, the eigenspaces of the pencil do not span the full space \mathbb{k}^n . However, the space \mathbb{k}^n admits a direct sum decomposition in terms of *root subspaces*, a generalization of the concept; see [Atk72, Theorem 6.9.2]. \diamond

4.2.2 Regularity of matrix pencils

In the following, we focus on sufficient conditions for a pencil of matrices to be regular. We start by showing that a pencil is regular as long as the matrices are sufficiently generic, which is the content of the following proposition.

Proposition 4.2.11. *Let \mathbb{k} be an infinite field and let $V_0, \dots, V_r \subseteq \mathbb{k}^n$ be vector subspaces. Let $\Delta_0, \dots, \Delta_r \in \mathbb{k}^{n \times n}$ be matrices that are generic among the matrices satisfying $\Delta_j V_j = 0$ for all $0 \leq j \leq r$. Then the pencil $(\Delta_0, \dots, \Delta_r)$ is regular if and only if $V_0 \cap \dots \cap V_r = 0$.*

Here, *generic* means that the property holds for all matrix pencils from some non-empty Zariski-open subset (which, in particular, is Zariski-dense since \mathbb{k} is infinite) of the space of all matrix pencils $(\Delta_0, \dots, \Delta_r)$ that satisfy $\Delta_j V_j = 0$ for $0 \leq j \leq r$. As the Zariski topology is rather coarse, this is a very weak assumption.

Also note that $\ker \Delta_j = V_j$ for all $0 \leq j \leq r$, since the matrices $\Delta_0, \dots, \Delta_r$ are generic, so another way to phrase Proposition 4.2.11 is that a generic such pencil $(\Delta_0, \dots, \Delta_r)$ is regular if and only if $\bigcap_{j=0}^r \ker \Delta_j = 0$.

Proof. If the pencil $(\Delta_0, \dots, \Delta_r)$ is regular, then clearly the spaces V_0, \dots, V_r must intersect trivially. This even holds for arbitrary (non-generic) matrices that satisfy $\Delta_j V_j = 0$, $0 \leq j \leq r$.

For the converse, assume that $V_0 \cap \dots \cap V_r = 0$. This implies that

$$V_0^\perp + \dots + V_r^\perp = (V_0 \cap \dots \cap V_r)^\perp = \mathbb{k}^n. \quad (4.5)$$

We will show that the matrix $\Delta_0 + \dots + \Delta_r$ is generically non-singular.

Denote by $W_j \subseteq \mathbb{k}^{n \times n}$ the subspaces $W_j = \{M \in \mathbb{k}^{n \times n} \mid M V_j = 0\}$ for $0 \leq j \leq r$. Note that the rows of a matrix $M \in W_j$ are vectors in V_j^\perp . We claim that the map

$$\begin{aligned} W_0 \times \dots \times W_r &\longrightarrow \mathbb{k}^{n \times n}, \\ (M_0, \dots, M_r) &\longmapsto M_0 + \dots + M_r, \end{aligned} \quad (4.6)$$

is surjective. As the property $M_j V_j = 0$ for $M_j \in W_j$ imposes linear relations between the entries of each row of M_j , it is enough to check the surjectivity on each row individually.

Let $M \in \mathbb{k}^{n \times n}$ be an arbitrary matrix and denote the rows of M by $u_1, \dots, u_n \in \mathbb{k}^n$. Due to (4.5), for every $1 \leq i \leq n$, we can choose vectors $u_i^{(j)} \in V_j^\perp$, $0 \leq j \leq r$, such that $\sum_{j=0}^r u_i^{(j)} = u_i$. Thus, by setting $M_j := (u_1^{(j)}, \dots, u_n^{(j)})^\top$, we have $M_j \in W_j$ and $\sum_{j=0}^r M_j = M$. This proves the claim that the map (4.6) is surjective.

The matrix $\Delta_0 + \dots + \Delta_r \in \mathbb{k}^{n \times n}$ is singular if and only if $\det(\Delta_0 + \dots + \Delta_r) = 0$. This determinant is a polynomial in the entries of $\Delta_0, \dots, \Delta_r$ which does not vanish on all of $W_0 \times \dots \times W_r$, since the image of the surjective map (4.6) contains non-singular matrices. Therefore, the set

$$\{(M_0, \dots, M_r) \in W_0 \times \dots \times W_r \mid \det(M_0 + \dots + M_r) \neq 0\}$$

is a non-empty Zariski-open subset of $W_0 \times \dots \times W_r$. If $(\Delta_0, \dots, \Delta_r)$ is contained in this set, then $\Delta_0 + \dots + \Delta_r$ is non-singular, so the pencil $(\Delta_0, \dots, \Delta_r)$ must be regular. \square

Example 4.2.12. If the matrix pencil $(\Delta_0, \dots, \Delta_r)$ in Proposition 4.2.11 is not generic, then the statement of the proposition does not hold in general. This means that the assumption $\bigcap_{j=0}^r \ker \Delta_j = 0$ is not enough to conclude that an arbitrary matrix pencil $(\Delta_0, \dots, \Delta_r)$ is regular, even though the converse is true. For instance, this becomes apparent by observing that the left kernels of $\Delta_0, \dots, \Delta_r$ might still have a non-trivial intersection, in which case the determinant of any linear combination of $\Delta_0, \dots, \Delta_r$ would be zero. Thus, Proposition 4.2.11 implies in particular that the left kernels of generic matrices as in the proposition must have a trivial intersection.

However, even if the left kernels intersect trivially as well, a non-generic pencil can still be singular, as the following example of a pencil of two matrices from [Sle+00, Chapter 8.7.4] shows. Define the matrices

$$\Delta_0 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \quad \Delta_1 = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Then the right kernels as well as the left kernels of Δ_0, Δ_1 each have trivial intersection. However, $\det(\lambda_0 \Delta_0 + \lambda_1 \Delta_1) = 0$ for every choice of λ_0, λ_1 , so the pencil (Δ_0, Δ_1) is singular. Note that the right kernel of $\lambda_0 \Delta_0 + \lambda_1 \Delta_1$ is spanned by the vector $(-\lambda_1, \lambda_0, 0)$, which varies for different choices of λ_0, λ_1 .

For a singular pencil of two matrices, this can also be read off of the *Kronecker Canonical Form* of the pencil; see [Sle+00, Chapter 8.7.2]. In terms of this formalism, the common right and left kernels correspond to the blocks L_0 and L_0^\top , respectively, so the singular structure contained in any of the blocks L_j, L_j^\top for $j > 0$ is still present after removing common kernels if any such blocks exist in the Kronecker Canonical Form of the pencil. \diamond

The following lemma shows that we can drop the genericity assumption of Proposition 4.2.11 if the matrices are positive-semidefinite.

Lemma 4.2.13. *Let $H_0, \dots, H_r \in \mathbb{C}^{n \times n}$ be positive-semidefinite matrices. Then the matrix pencil (H_0, \dots, H_r) is non-singular if and only if $\bigcap_{j=0}^r \ker H_j = 0$.*

Proof. If the pencil is non-singular, there exists an invertible linear combination of the matrices H_0, \dots, H_r , so the kernels must intersect trivially. Conversely, assume that $\bigcap_{j=0}^r \ker H_j = 0$. We claim that any linear combination $H = \sum_{j=0}^r \lambda_j H_j$ with $\lambda_0, \dots, \lambda_r > 0$ is non-singular. Indeed, if $p \in \ker H$, then

$$p^* H p = \sum_{j=0}^r \lambda_j p^* H_j p = 0.$$

Since $p^* H_j p \geq 0$ and $\lambda_j > 0$, this implies that in fact $p \in \ker H_j$ for $0 \leq j \leq r$ and thus $p = 0$. \square

The positive-semidefiniteness of the matrices in Lemma 4.2.13 is important for the statement. In Example 4.2.12, we have seen a counter-example involving matrices that are not positive-semidefinite.

We end this section by looking more closely at an example for Lemma 4.2.13. It also serves as motivation for the study of eigenvalues and eigenspaces of moment matrices associated to measures supported on algebraic varieties, which we further pursue in the rest of the chapter.

Example 4.2.14. Let \mathcal{N} be the standard Gaussian distribution on \mathbb{R} , i.e. the measure with density $\frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}x^2}$. We consider the measures $\mu_1 := \mathcal{N} \otimes \delta_0$ and $\mu_2 := \delta_0 \otimes \mathcal{N}$ on the product space \mathbb{R}^2 , where δ_0 denotes the Dirac measure. Then the coordinate axes in \mathbb{R}^2 are the joint support of the measures μ_1, μ_2 . This is an algebraic variety defined by the ideal $\langle x_1x_2 \rangle \subseteq R := \mathbb{R}[x_1, x_2] \cong \mathbb{R}[x] \otimes \mathbb{R}[x]$.

For $i = 1, 2$, let $H_i := \left(\int_{\mathbb{R}^2} x^{\alpha+\beta} d\mu_i(x) \right)_{|\alpha|, |\beta| \leq 2}$ be the moment matrix truncated at degree 2 corresponding to the measure μ_i . Due to symmetry, the odd moments vanish, which we indicate by leaving the respective entries of the moment matrices empty. With respect to the basis $1, x_1, x_2, x_1^2, x_1x_2, x_2^2$, the moment matrices are

$$H_1 = \left(\begin{array}{c|cc|ccc} 1 & & & 0 & 0 & 1 \\ \hline & 0 & 0 & & & \\ & 0 & 1 & & & \\ \hline 0 & & & 0 & 0 & 0 \\ 0 & & & 0 & 0 & 0 \\ 1 & & & 0 & 0 & 3 \end{array} \right), \quad H_2 = \left(\begin{array}{c|cc|ccc} 1 & & & 1 & 0 & 0 \\ \hline & 1 & 0 & & & \\ & 0 & 0 & & & \\ \hline 1 & & & 3 & 0 & 0 \\ 0 & & & 0 & 0 & 0 \\ 0 & & & 0 & 0 & 0 \end{array} \right).$$

These are positive-semidefinite matrices of rank 3 that form a singular pencil. The common kernel of H_1 and H_2 is one-dimensional and is spanned by x_1x_2 as a subspace of $R_{\leq 2}$. In this example, by removing the respective row and column consisting of zeros, we obtain a regular pencil by Lemma 4.2.13. The eigenvalues of the regular pencil are $[0 : 1]$, $[1 : 0]$ (each with multiplicity 2) as well as $[1 : 1]$.

As a subspace of $R_{\leq 2}/(\langle x_1x_2 \rangle \cap R_{\leq 2})$, the eigenspace corresponding to the eigenvalue $[0 : 1]$ is spanned by the polynomials x_1, x_1^2 which vanish on $\text{supp } \mu_2$, the x_2 -axis. Similarly, the eigenspace for $[1 : 0]$ is spanned by x_2, x_2^2 , vanishing on the other axis. The eigenspace corresponding to $[1 : 1]$ is spanned by $1 - \frac{1}{3}(x_1^2 + x_2^2)$ which is a polynomial that does not vanish on either of the two axes.

In summary, we observe that some of the eigenvectors are polynomials that vanish on some component of the underlying algebraic variety, while for other eigenvectors this is not the case. We will examine this property more closely in Section 4.5 and use it to recover the vanishing ideals of each individual component by solving a generalized eigenvalue problem. \diamond

4.2.3 Orthogonality

Given an eigenvector p of a matrix pencil (H_0, \dots, H_r) , we collect a few criteria for a non-zero vector q to be orthogonal to all the vectors H_0p, \dots, H_rp .

Lemma 4.2.15. *Let (H_0, \dots, H_r) be a regular matrix pencil over \mathbb{k} and let H be a non-singular linear combination of H_0, \dots, H_r . Let p be an eigenvector of the pencil and let q be any vector. Then $q^\top H p = 0$ if and only if $q^\top H_j p = 0$ for all $0 \leq j \leq r$.*

Proof. As p is an eigenvector of the pencil, we have

$$H_j p = \gamma_j H p$$

for $0 \leq j \leq r$ and a suitable choice of the coordinates $\gamma_0, \dots, \gamma_r \in \mathbb{k}$. Hence, if $q^\top H p = 0$, then also $q^\top H_j p = \gamma_j q^\top H p = 0$ for all $0 \leq j \leq r$. The converse follows immediately from the fact that H is a linear combination of H_0, \dots, H_r . \square

Lemma 4.2.16. *Let (H_0, \dots, H_r) be a regular pencil of complex matrices. Let H be a linear combination of H_0, \dots, H_r which is an invertible matrix. Let p be an eigenvector of (H_0, \dots, H_r) with eigenvalue $[\gamma] \in \mathbb{P}_{\mathbb{C}}^r$ for some $\gamma \in \mathbb{C}^{r+1}$ such that $H_j p = \gamma_j H p$ for all $0 \leq j \leq r$. Further, let $\lambda \in \mathbb{C}$ and let q be an eigenvector of the matrix $(H^*)^{-1} H_0^*$ with eigenvalue $\bar{\lambda}$. If $\lambda \neq \gamma_0$, then $q^* H_j p = 0$ for all $0 \leq j \leq r$.*

Proof. Since q is an eigenvector with eigenvalue $\bar{\lambda}$ of the matrix $(H^*)^{-1} H_0^*$, we have $H_0^* q = \bar{\lambda} H^* q$ or, equivalently, $q^* H_0 = \lambda q^* H$. Then it follows that

$$\gamma_0 q^* H p = q^* H_0 p = \lambda q^* H p.$$

As $\gamma_0 \neq \lambda$, this implies that $q^* H p = 0$. Then the conclusion follows from Lemma 4.2.15, since p is an eigenvector of the pencil (H_0, \dots, H_r) . \square

Corollary 4.2.17. *Let (H_0, \dots, H_r) be a regular pencil of complex matrices and let p be an eigenvector of the pencil with eigenvalue $\gamma \in \mathbb{P}_{\mathbb{C}}^r$. Then the following properties hold:*

- (1) *If q is an eigenvector of (H_0^*, \dots, H_r^*) with eigenvalue $\gamma' \in \mathbb{P}_{\mathbb{C}}^r$ such that $\gamma' \neq \bar{\gamma}$, then $q^* H_j p = 0$ for all $0 \leq j \leq r$.*
- (2) *If q is an eigenvector of $(H_0^\top, \dots, H_r^\top)$ with eigenvalue $\gamma' \in \mathbb{P}_{\mathbb{C}}^r$ such that $\gamma' \neq \gamma$, then $q^\top H_j p = 0$ for all $0 \leq j \leq r$.*
- (3) *If the matrices H_0, \dots, H_r are Hermitian and $\gamma \neq \bar{\gamma}$, then $p^* H_j p = 0$ for all $0 \leq j \leq r$.*
- (4) *If $q \in \ker H_0^*$ and $\gamma_0 \neq 0$, then $q^* H_j p = 0$ for all $0 \leq j \leq r$.*

Proof. For (1), we can choose representatives of γ and γ' such that $H_j p = \gamma_j H p$ and $H_j^* q = \gamma'_j H^* q$ for all $0 \leq j \leq r$, where H denotes some non-singular linear combination of H_0, \dots, H_r . As $\gamma \neq \bar{\gamma}'$, we can assume, without loss of generality, that $\gamma_0 \neq \bar{\gamma}'_0$, so the result follows from Lemma 4.2.16, since $(H^*)^{-1} H_0^* q = \gamma'_0 q$.

Part (2) follows directly from (1) by observing that \bar{q} is an eigenvector with eigenvalue $\bar{\gamma}'$ of the pencil (H_0^*, \dots, H_r^*) and part (3) is a special case of (1).

Finally, for (4), without loss of generality we can assume that $q \neq 0$. Let H be a non-singular linear combination of H_0, \dots, H_r as before. Since $q \in \ker H_0^*$, we have that q is

an eigenvector of $(H^*)^{-1}H_0^*$ for the eigenvalue 0. As $\gamma_0 \neq 0$, the result then follows from Lemma 4.2.16. \square

Remark 4.2.18. Particularly important for us is the following consequence of Corollary 4.2.17(1) which holds by contraposition. Assume that (H_0, \dots, H_r) is a regular pencil of Hermitian matrices and that p, q are eigenvectors with distinct eigenvalues $\gamma, \gamma' \in \mathbb{P}_{\mathbb{C}}^r$, $\gamma \neq \gamma'$. If $q^*H_j p \neq 0$ for some $j \in \{0, \dots, r\}$, then it follows that $\gamma = \overline{\gamma'}$. In particular, this implies that γ, γ' cannot have real coordinates. See Lemma 4.2.22 for an extension of this argument. Also note that the converse of Corollary 4.2.17(3) is not true in general. \diamond

Example 4.2.19. Let $a \in \mathbb{R}_{>0}$ and define the Hermitian matrices

$$H_0 = \begin{pmatrix} -a^2 & i \\ -i & 1 \end{pmatrix}, \quad H_1 = \begin{pmatrix} -a^2 & 0 \\ 0 & 1 \end{pmatrix}.$$

Then (H_0, H_1) is a regular pencil with eigenvalues $\gamma, \gamma' \in \mathbb{P}_{\mathbb{C}}^1$ and respective eigenvectors p, q where

$$\gamma = [a - i : a], \quad \gamma' = [a + i : a], \quad p = \begin{pmatrix} 1 \\ a \end{pmatrix}, \quad q = \begin{pmatrix} 1 \\ -a \end{pmatrix}.$$

From this, we obtain

$$q^*H_0p = -2a(a - i) \neq 0 \quad \text{and} \quad q^*H_1p = -2a^2 \neq 0,$$

so the statement of Corollary 4.2.17(1) can fail if $\gamma = \overline{\gamma'}$. \diamond

4.2.4 Linear transformations of generalized eigenproblems

The following result plays an important role for the algorithms we develop later on. It describes what happens to a regular eigenvalue problem under a linear transformation.

Proposition 4.2.20. *Let $H_0, \dots, H_r \in \mathbb{k}^{n \times n}$, $r \geq 1$, be matrices such that the pencil (H_0, \dots, H_r) is regular. Let $s \geq r$ and define matrices $M_i := \sum_{j=0}^r \lambda_{ij} H_j$ for $0 \leq i \leq s$, where $(\lambda_{ij})_{ij} \in \mathbb{k}^{(s+1) \times (r+1)}$ is a matrix of rank $r+1$. Then the following properties hold:*

- (1) *The pencil (M_0, \dots, M_s) is regular.*
- (2) *A vector $p \neq 0$ is an eigenvector of (H_0, \dots, H_r) if and only if it is an eigenvector of (M_0, \dots, M_s) .*
- (3) *The matrix $(\lambda_{ij})_{ij}$ induces a 1-to-1 correspondence between the eigenvalues in $\mathbb{P}_{\mathbb{k}}^r$ of (H_0, \dots, H_r) and the eigenvalues in $\mathbb{P}_{\mathbb{k}}^s$ of (M_0, \dots, M_s) .*

Proof. As (H_0, \dots, H_r) is a regular pencil, we can choose a linear combination $H := \sum_{j=0}^r c_j H_j$ which is invertible for suitable $c_j \in \mathbb{k}$. Moreover, we can assume that $H_j = \sum_{i=0}^s \mu_{ji} M_i$ for a suitable choice of $(\mu_{ji})_{ji} \in \mathbb{k}^{(r+1) \times (s+1)}$, since the matrix $(\lambda_{ij})_{ij}$ has full column rank. Then

$$H = \sum_{j=0}^r c_j H_j = \sum_{i=0}^s \left(\sum_{j=0}^r c_j \mu_{ji} \right) M_i$$

is an invertible linear combination of M_0, \dots, M_s , so the pencil (M_0, \dots, M_s) is regular.

If $p \neq 0$ is an eigenvector of (H_0, \dots, H_r) with eigenvalue $\gamma \in \mathbb{P}^r$ such that $H_j p = \gamma_j H p$ for all $0 \leq j \leq r$, then

$$M_i p = \sum_{j=0}^r \lambda_{ij} H_j p = \left(\sum_{j=0}^r \lambda_{ij} \gamma_j \right) H p, \quad (4.7)$$

so p is an eigenvector of (M_0, \dots, M_s) with eigenvalue $[\sum_{j=0}^r \lambda_{ij} \gamma_j]_{0 \leq i \leq s} \in \mathbb{P}_{\mathbb{k}}^s$, which is well-defined as $(\lambda_{ij})_{ij}$ has full rank. Conversely, if $p \neq 0$ is an eigenvector of (M_0, \dots, M_s) with eigenvalue $\delta \in \mathbb{P}^s$ such that $M_i p = \delta_i H p$ for all $0 \leq i \leq s$, then

$$H_j p = \sum_{i=0}^s \mu_{ji} M_i p = \left(\sum_{i=0}^s \mu_{ji} \delta_i \right) H p. \quad (4.8)$$

If $\sum_{i=0}^s \mu_{ji} \delta_i = 0$ for all $0 \leq j \leq r$, we obtain $H_j p = 0$ and consequently $H p = 0$. As H is invertible, this is a contradiction to the assumption that $p \neq 0$. Therefore, $\sum_{i=0}^s \mu_{ji} \delta_i$ is non-zero for some j , so p is an eigenvector of (H_0, \dots, H_r) with eigenvalue $[\sum_{i=0}^s \mu_{ji} \delta_i]_{0 \leq j \leq r} \in \mathbb{P}_{\mathbb{k}}^r$.

It remains to show that $(\lambda_{ij})_{ij}$ induces a 1-to-1 correspondence between the eigenvalues of the two pencils. Denote by $\Gamma \subset \mathbb{P}_{\mathbb{k}}^r$ and $\Gamma' \subseteq \mathbb{P}_{\mathbb{k}}^s$ the set of eigenvalues of (H_0, \dots, H_r) and (M_0, \dots, M_s) , respectively. The computation above implies that $(\lambda_{ij})_{ij}$ and $(\mu_{ji})_{ji}$ induce well-defined maps $\iota: \Gamma \rightarrow \Gamma'$ and $\tau: \Gamma' \rightarrow \Gamma$. The map ι is injective, but for τ this is not clear when $s > r$. However, if p is an eigenvector of (M_0, \dots, M_s) with eigenvalue $\delta \in \mathbb{P}_{\mathbb{k}}^s$, then, as shown above, p is also an eigenvector of (M_0, \dots, M_s) with eigenvalue $\iota(\tau(\delta)) \in \mathbb{P}_{\mathbb{k}}^s$. As the eigenvalue for p in $\mathbb{P}_{\mathbb{k}}^s$ is unique by Lemma 4.2.8, this means that τ must be injective as well, so it is the inverse of ι . \square

Remark 4.2.21. On the level of coordinates of the eigenvalues, the proof of Proposition 4.2.20 shows that the injective map $\Lambda: \mathbb{k}^{r+1} \rightarrow \mathbb{k}^{s+1}$ defined by $(\lambda_{ij})_{ij}$ gives rise to a 1-to-1 correspondence between representatives in $\mathbb{k}^{r+1} \setminus \{0\}$ and $\mathbb{k}^{s+1} \setminus \{0\}$ of eigenvalues of (H_0, \dots, H_r) and (M_0, \dots, M_s) , respectively, as long as the coordinates are chosen with respect to the fixed matrix H . As described in Corollary 4.2.7, this choice is uniquely determined. Thus, if in those coordinates $\gamma \in \mathbb{k}^{r+1} \setminus \{0\}$ and $\delta = \Lambda(\gamma) \in \mathbb{k}^{s+1} \setminus \{0\}$ are representatives of eigenvalues of the two pencils for some eigenvector p , then these coordinates satisfy $H_j p = \gamma_j H p$ and $M_i p = \delta_i H p$ for all $0 \leq j \leq r$ and $0 \leq i \leq s$, as follows from Equations (4.7) and (4.8). In other words, the map Λ is compatible with this choice of coordinates. \diamond

Lemma 4.2.22. *Let (M_0, \dots, M_s) be a regular pencil of complex matrices of the form $M_i = \sum_{j=0}^r \lambda_{ij} H_j$, $0 \leq i \leq s$, where H_j , $0 \leq j \leq r$, are Hermitian matrices and $(\lambda_{ij})_{ij} \in \mathbb{C}^{(s+1) \times (r+1)}$ with $s \geq r$ is a matrix of rank $r+1$. Denote by $\Lambda: \mathbb{P}_{\mathbb{C}}^r \rightarrow \mathbb{P}_{\mathbb{C}}^s$ the map induced by $(\lambda_{ij})_{ij}$. Let M be a non-singular matrix that is a linear combination of M_0, \dots, M_s . If p, q are eigenvectors of the pencil (M_0, \dots, M_s) with eigenvalues $\gamma, \gamma' \in \mathbb{P}_{\mathbb{C}}^s$ with $\gamma \neq \gamma'$ that satisfy $q^* M p \neq 0$, then $\Lambda^{-1}(\gamma) = \overline{\Lambda^{-1}(\gamma')}$. In particular,*

$$\Lambda^{-1}(\gamma), \Lambda^{-1}(\gamma') \notin \mathbb{P}_{\mathbb{R}}^r.$$

Proof. Observe that (H_0, \dots, H_r) is a regular matrix pencil, since (M_0, \dots, M_s) is regular, so M is also a linear combination of H_0, \dots, H_r . As the matrix $(\lambda_{ij})_{ij}$ has full rank, the map Λ is injective and induces a 1-to-1 correspondence between the eigenvalues of the pencils by Proposition 4.2.20.

As M is a non-singular linear combination of H_0, \dots, H_r and since $q^* M p \neq 0$, by Lemma 4.2.15, we can assume that $q^* H_j p \neq 0$ for some $j \in \{0, \dots, r\}$. Then, since the matrices H_0, \dots, H_r are Hermitian and p, q are eigenvectors of (H_0, \dots, H_r) with eigenvalues $\Lambda^{-1}(\gamma), \Lambda^{-1}(\gamma')$, it follows from Corollary 4.2.17 (1) that $\Lambda^{-1}(\gamma) = \overline{\Lambda^{-1}(\gamma')}$. However, since $\gamma \neq \gamma'$, this implies that $\Lambda^{-1}(\gamma), \Lambda^{-1}(\gamma')$ are distinct and cannot have real coordinates. \square

4.3 Hermitian and symmetric cases

In this section, we derive a first recovery algorithm, Algorithm 4.1, for the case of Hermitian or complex symmetric matrices. It does not assume that the matrices are moment matrices, but is phrased more generally for any such matrices. The algorithm is primarily expository in nature, serving as foundation for the variants that we develop later in Sections 4.4 and 4.7. We demonstrate the use of Algorithm 4.1 as well as its shortcomings by zero- and positive-dimensional examples in Sections 4.3.1 and 4.3.2.

For clarity, we start by mentioning the following simple equivalence.

Lemma 4.3.1. *Let $U, U' \subseteq V$ be submodules of a module V over a commutative ring. Then $U \setminus U' \neq \emptyset$ if and only if $U/(U' \cap U) \neq 0$.*

Proof. The equality $U/(U' \cap U) = 0$ holds if and only if $U = U' \cap U$ or, equivalently, $U \subseteq U'$, which is equivalent to $U \setminus U' = \emptyset$. \square

Remark 4.3.2. Thus, for a pencil (H_0, \dots, H_r) and all $0 \leq j \leq r$, we have that $\bigcap_{k=0, k \neq j}^r \ker H_k \setminus \ker H_j \neq \emptyset$ if and only if $\bigcap_{k=0, k \neq j}^r \ker H_k / (\bigcap_{k=0}^r \ker H_k) \neq 0$. This is a property that will frequently play a role, in the remainder of the chapter, as it is a prerequisite for the eigenvalue-based recovery approach. If additionally the pencil is regular, the above is equivalent to $\bigcap_{k=0, k \neq j}^r \ker H_k \neq 0$. To see the latter, let $p \in \bigcap_{k=0, k \neq j}^r \ker H_k$,

$p \neq 0$, and let H be a non-singular linear combination of H_0, \dots, H_r . Then the assumption $p \in \ker H_j$ implies that $Hp = 0$, which is a contradiction to the hypothesis that $p \neq 0$. Hence, we must have $p \notin \ker H_j$. \diamond

Remark 4.3.3. In later sections, we will mainly focus on situations in which the kernels $\ker H_j$ of the matrices can be viewed as certain subspaces of ideals, say $\mathfrak{a}_j \cap U$, where $U \subseteq R$ is a vector subspace and $\mathfrak{a}_j \subseteq R$ is an ideal in a polynomial ring $R = \mathbb{k}[x_1, \dots, x_n]$, for $0 \leq j \leq r$.

Recall that, if $n = 1$, a *Lagrange polynomial* associated to a set of distinct points $\xi_0, \dots, \xi_r \in \mathbb{k}$ is a polynomial that vanishes on exactly r of these points. This concept is also meaningful in higher dimensions for points $\xi_0, \dots, \xi_r \in \mathbb{k}^n$, $n \geq 1$ (cf. [Isk18, Chapter 8.1]).

From this point of view, we can interpret a polynomial $p \in \bigcap_{k=0, k \neq j}^r \ker H_k \setminus \ker H_j = \bigcap_{k=0, k \neq j}^r (\mathfrak{a}_k \cap U) \setminus \mathfrak{a}_j \neq \emptyset$ as a generalization of a Lagrange polynomial. Indeed, if the ideals are zero-dimensional of the form $\mathfrak{a}_k = \langle x_1 - \xi_{k1}, \dots, x_n - \xi_{kn} \rangle$, $0 \leq k \leq r$, then p vanishes at $\{\xi_0, \dots, \xi_r\} \setminus \{\xi_j\}$. More generally, for arbitrary ideals, p vanishes on the varieties $V(\mathfrak{a}_k)$ with $0 \leq k \leq r$, $k \neq j$. (In particular, some of the eigenvectors in Example 4.2.14 are polynomials of this kind.) Note however that p may vanish at some points of $V(\mathfrak{a}_j)$, as well. In fact, depending on the dimensions of the varieties, this is not uncommon as the varieties can have intersection points. Though, if \mathbb{k} is algebraically closed and \mathfrak{a}_j is a radical ideal, then p does not fully vanish on $V(\mathfrak{a}_j)$ by construction, since $p \notin \mathfrak{a}_j = \sqrt{\mathfrak{a}_j} = I(V(\mathfrak{a}_j))$. \diamond

Example 4.3.4. We consider a minimal example of two zero-dimensional components, that is, two points $\xi_1 = (2, 1), \xi_2 = (-5, 3) \in \mathbb{k}^2$. Let $\mathfrak{m}_{\xi_j} = \langle x - \xi_j \rangle \subseteq R := \mathbb{k}[x_1, x_2]$, $j = 1, 2$, be the corresponding point ideals. Then $R/(\mathfrak{m}_{\xi_1} \cap \mathfrak{m}_{\xi_2})$ is a two-dimensional vector space and we can choose $1, x_2 - 2 \in R$ as representatives of a vector space basis.

Let $\sigma_1 = \frac{1}{2} \text{ev}_{\xi_1}, \sigma_2 = \frac{1}{2} \text{ev}_{\xi_2}$. Then, in terms of the chosen basis, we can write the Gramian matrix corresponding to the symmetric bilinear form $\langle -, - \rangle_{\sigma_1 + \sigma_2}$, defined in Definition 3.1.6, as

$$M_0 = (\sigma_1 + \sigma_2) \begin{pmatrix} 1 & x_2 - 2 \\ x_2 - 2 & (x_2 - 2)^2 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

In other words, the basis is orthonormal with respect to $\langle -, - \rangle_{\sigma_1 + \sigma_2}$. For a weighted functional such as $\tau := 3\sigma_1 + 4\sigma_2$, we similarly obtain a matrix

$$M_1 = \tau \begin{pmatrix} 1 & x_2 - 2 \\ x_2 - 2 & (x_2 - 2)^2 \end{pmatrix} = \begin{pmatrix} \frac{7}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{7}{2} \end{pmatrix}.$$

Since M_0 is the identity matrix, in order to obtain the eigenvalues of the pencil (M_0, M_1) , it is enough to compute the eigenvalues of M_1 , which are 3 and 4, so they agree with the weights of τ . Thus, the pencil (M_0, M_1) has eigenvalues $[1 : 3], [1 : 4] \in \mathbb{P}_{\mathbb{k}}^1$. The corresponding eigenvectors can be chosen as $(1, -1) = x_2 - 3$ and $(1, 1) = x_2 - 1$,

respectively. These polynomials vanish on the opposite component, so the eigenvector for eigenvalue 3 vanishes on ξ_2 and the one for 4 on ξ_1 . Hence, they are Lagrange-like polynomials as explained in Remark 4.3.3. It will turn out that this is not a coincidence and we will make further use of these observations in Section 4.5. \diamond

Algorithm 4.1 Recovery of components from several weighted sums

Input: A natural number $s \in \mathbb{N}$ and a regular pencil (M_0, \dots, M_s) of complex matrices.

Assumptions: The matrices are of the form $M_i = \sum_{j=0}^r \lambda_{ij} H_j$, $0 \leq i \leq s$, for some $r \leq s$, where $H_0, \dots, H_r \in \mathbb{C}^{m \times m}$ are Hermitian (or complex symmetric) matrices and where $(\lambda_{ij})_{ij} \in \mathbb{C}^{(s+1) \times (r+1)}$ is a matrix of rank $r+1$. Furthermore, we have $\bigcap_{k=0, k \neq j}^r \ker H_k \neq 0$, for every $0 \leq j \leq r$.

Output: r as well as $[\lambda_{0j} : \dots : \lambda_{sj}] \in \mathbb{P}_{\mathbb{C}}^s$ and $[H_j] \in \mathbb{P}(\mathbb{C}^{m \times m})$ for all $0 \leq j \leq r$ (up to scaling and permutation); or failure.

- 1: Let M be a linear combination of M_0, \dots, M_s that is a non-singular matrix.
- 2: Compute the set \mathfrak{V} of eigenspaces of the matrix pencil (M_0, \dots, M_s) . Denote by $\Gamma \subseteq \mathbb{P}_{\mathbb{C}}^s$ the set of corresponding eigenvalues.
- 3: Let r be the dimension of the projective subspace spanned by Γ .
- 4: If Γ consists of exactly $r+1$ eigenvalues, go to Line 9.
- 5: **if** in the Hermitian case **then**
- 6: **for** every pair $\gamma, \gamma' \subseteq \Gamma$, $\gamma \neq \gamma'$, with corresponding eigenspaces $V, V' \in \mathfrak{V}$ **do**
- 7: If V' is not orthogonal to MV , remove γ and γ' from Γ .
- 8: If Γ consists of more than $r+1$ eigenvalues, fail.
- 9: With $\Gamma = \{\gamma_0, \dots, \gamma_r\}$, we have $[\lambda_{0j} : \dots : \lambda_{sj}] = \gamma_j$ for $0 \leq j \leq r$ (up to permutations in j).
- 10: Let $(\tilde{\lambda}_{0j}, \dots, \tilde{\lambda}_{sj})$ be a representative of $[\lambda_{0j} : \dots : \lambda_{sj}]$ for $0 \leq j \leq r$ and solve the linear system

$$\sum_{j=0}^r \tilde{\lambda}_{ij} \tilde{H}_j = M_i, \quad 0 \leq i \leq s,$$

for $\tilde{H}_0, \dots, \tilde{H}_r$. Then \tilde{H}_j is a representative of $[H_j] \in \mathbb{P}(\mathbb{C}^{m \times m})$ for $0 \leq j \leq r$.

Theorem 4.3.5. *Algorithm 4.1 works unless the check in Line 8 fails.*

Proof. First observe that (H_0, \dots, H_r) is a regular matrix pencil, since (M_0, \dots, M_s) is regular. As the matrix $(\lambda_{ij})_{ij}$ has full rank, then, by Proposition 4.2.20, the eigenvectors of the two pencils agree and there is a 1-to-1 correspondence between their eigenvalues.

By assumption, the subspaces $U_j := \bigcap_{k=0, k \neq j}^r \ker H_k$ are non-trivial for every $0 \leq j \leq r$. Let $j \in \{0, \dots, r\}$ be arbitrary. If $p \in U_j$, $p \neq 0$, then $p \notin \ker H_j$, as explained in Remark 4.3.2, since the pencil (H_0, \dots, H_r) is regular. Thus, p is an eigenvector of the pencil (H_0, \dots, H_r) with eigenvalue $e_j := [\delta_{kj}]_{0 \leq k \leq r} = [0 : \dots : 0 : 1 : 0 : \dots : 0] \in \mathbb{P}_{\mathbb{C}}^r$, since

$$H_k p = 0 \cdot H_j p = 0$$

for all $0 \leq k \leq r$ with $k \neq j$. As j was arbitrary, this means that e_0, \dots, e_r arise as

eigenvalues of the pencil (H_0, \dots, H_r) .

Denote by $\Lambda: \mathbb{P}_{\mathbb{C}}^r \rightarrow \mathbb{P}_{\mathbb{C}}^s$ the injective map induced by $(\lambda_{ij})_{ij}$. By Proposition 4.2.20 (3), it then follows that $[\lambda_{0j} : \dots : \lambda_{sj}] = \Lambda(e_j)$, $0 \leq j \leq r$, are eigenvalues of (M_0, \dots, M_s) , so they are contained in the set Γ computed in Line 2. As every eigenvalue in Γ must be contained in the image of the injective map Λ , we obtain r as the dimension of the projective subspace spanned by Γ , in Line 3. In particular, Γ has at least cardinality $r + 1$ in Line 4.

As we are only interested in those eigenvalues $[\lambda_{0j} : \dots : \lambda_{sj}]$ in Γ which correspond to the eigenvalues e_j , $0 \leq j \leq r$, of the pencil (H_0, \dots, H_r) , we can filter out some unwanted eigenvalues by the check in Line 7, in the Hermitian case. This works because of Lemma 4.2.22, as $e_j = \Lambda^{-1}([\lambda_{0j} : \dots : \lambda_{sj}])$ has real coordinates. (In the symmetric case, the corresponding check is redundant by Corollary 4.2.17 (2), so we skip it.)

We only remove pairs of eigenvalues γ, γ' from Γ where $\Lambda^{-1}(\gamma) = \overline{\Lambda^{-1}(\gamma')}$ and which do not allow for real preimages. Therefore, removing these as early as possible in the loop is not a problem by Corollary 4.2.17 (1), since for every other eigenvalue $\gamma'' \in \Gamma$ we must have $\Lambda^{-1}(\gamma'') \neq \overline{\Lambda^{-1}(\gamma)}, \overline{\Lambda^{-1}(\gamma')}$.

The check in Line 8 can only fail if the cardinality of Γ is larger than $r + 1$; as $\Lambda^{-1}(\Gamma)$ still contains e_0, \dots, e_r , the cardinality is at least $r + 1$. If it succeeds, we thus find $[\lambda_{0j} : \dots : \lambda_{sj}]$, $0 \leq j \leq r$, in Line 9.

Finally, assume that $\tilde{\lambda}_{ij} = c_j \lambda_{ij}$, $0 \leq i \leq s$, for some non-zero scalars c_j , $0 \leq j \leq r$. Letting $\tilde{H}_j = c_j^{-1} H_j$, the linear system

$$\sum_{j=0}^r \tilde{\lambda}_{ij} \tilde{H}_j = \sum_{j=0}^r \lambda_{ij} H_j = M_i, \quad 0 \leq i \leq s,$$

in Line 10 can be solved uniquely for $\tilde{H}_0, \dots, \tilde{H}_r$, as the matrix $(\tilde{\lambda}_{ij})_{ij}$ is of full rank. \square

Remark 4.3.6. If additionally the matrices H_0, \dots, H_r in Algorithm 4.1 are positive-semidefinite Hermitian matrices, then all the eigenvalues of the pencil (H_0, \dots, H_r) can be chosen to be non-negative, as we show in Corollary 4.4.1. Therefore, the check in Line 7 is redundant and can be skipped in this case. Section 4.4 gives a refinement of the algorithm to this situation.

The algorithm may also be adapted to the case in which the matrices H_0, \dots, H_r , instead of being Hermitian or symmetric, are arbitrary complex matrices or matrices over another field. In this case, the check in Line 7 must then be skipped as it does not apply. See for instance Example 4.3.16. \diamond

Algorithm 4.1 can fail when there are additional eigenvalues that have real preimages under the map induced by $(\lambda_{ij})_{ij}$ that are not unit vectors. See Example 4.3.19 for a particular example in which this happens.

Remark 4.3.7. If Algorithm 4.1 fails in Line 8, a heuristic approach for picking the correct $r + 1$ eigenvalues might consist of choosing the $r + 1$ eigenspaces $V \in \mathfrak{V}$ of highest dimension. The rationale behind this is that the multiplicity of the eigenvalue e_j of the pencil (H_0, \dots, H_r) is at least as large as the dimension of $\bigcap_{k=0, k \neq j}^r \ker H_k / \bigcap_{k=0}^r \ker H_k$, which is sometimes known to be of higher dimension than 1.

For instance, if the matrices are moment matrices of measures supported on positive-dimensional algebraic varieties given by ideals $\mathfrak{a}_0, \dots, \mathfrak{a}_r$ of the polynomial ring R , then increasing the degree leads to higher-dimensional eigenspaces. Recall that the subspaces $\mathfrak{a}_j \cap R_{\leq d}$ are related to the kernels of moment matrices, by Theorem 3.4.11. Under the assumptions of Lemma 1.2.4(1), the dimension of the space

$$\left(\bigcap_{k=0, k \neq j}^r \mathfrak{a}_k \cap R_{\leq d} \right) / \left(\bigcap_{k=0}^r \mathfrak{a}_k \cap R_{\leq d} \right)$$

then indeed grows with $d \in \mathbb{N}$, for each $0 \leq j \leq r$. As explained in Remark 4.3.3, we may interpret elements of this space as Lagrange-like polynomials on the component $V(\mathfrak{a}_j)$, vanishing on all the other components.

In cases in which the kernels correspond to vanishing ideals as above, another heuristic could use the inclusion properties of eigenspaces for growing degree (cf. Example 4.3.14). In Section 4.5, we expand on this idea in order to formulate a sufficient condition for asserting that (a variant of) this algorithm does not fail. \diamond

Remark 4.3.8. If in Algorithm 4.1 the matrices H_0, \dots, H_r are Hermitian matrices of rank 1, then the kernels $\ker H_j$, $0 \leq j \leq r$, which are of codimension 1, can be computed from the set of eigenvectors of the pencil (M_0, \dots, M_s) . This can serve as an alternative to solving the linear system of equations in Line 10, which possibly is more stable numerically.

For this, let p be an eigenvector of the pencil of the form $p \in \bigcap_{j=1}^r \ker H_j \setminus \ker H_0$, for instance. Let q be an arbitrary vector from the orthogonal complement of Mp , so $q^*Mp = 0$. By Lemma 4.2.15, this implies that $q^*H_0p = 0$. As H_0 is of rank 1, the vector H_0p spans the column space of H_0 , so $q \in \ker H_0^* = \ker H_0$. Hence, the complement of Mp recovers $\ker H_0$ and, proceeding similarly with the other r eigenvectors, one obtains $\ker H_1, \dots, \ker H_r$. The case of symmetric rank-1 matrices follows by a similar argument. \diamond

4.3.1 Zero-dimensional examples

As first examples, we apply Algorithm 4.1 to weighted sums of measures supported on zero-dimensional varieties. First, we study the case in which each component is a single point, for which reconstruction succeeds by Proposition 4.3.11 below. Recall that the multivariate Prony method is applicable in this situation, which is summarized in Section 1.3.1 and is covered extensively in [vdOhe17; Mou18] – here we primarily focus

on the application of Algorithm 4.1. Next, in Example 4.3.14, we examine the case in which the individual components consist of distinct sets of several points.

The following lemma provides several equivalent characterizations for the crucial requirement for Prony's method (cf. Theorem 1.3.1) that a certain evaluation map is surjective. This characterization is represented by case (4). As such, Lemma 4.3.9 gives insight into potential interpretations in a positive-dimensional context of this important criterion. This is discussed in more detail in Remark 4.3.10.

Lemma 4.3.9. *Let $\xi_0, \dots, \xi_r \in \mathbb{C}^n$ be distinct points and let $H_j = \left(\xi_j^{\alpha+\beta} \right)_{|\alpha|, |\beta| \leq d}$, $0 \leq j \leq r$, for some $d \in \mathbb{N}$. Let $R = \mathbb{C}[x_1, \dots, x_n]$ and define the ideals $\mathfrak{m}_{\xi_j} := \langle x - \xi_j \rangle \subseteq R$, $0 \leq j \leq r$. Then the following properties are equivalent:*

- (1) $\bigcap_{k=0, k \neq j}^r \ker H_k / \bigcap_{k=0}^r \ker H_k \neq 0$ for all $0 \leq j \leq r$;
- (2) $\bigcap_{k=0, k \neq j}^r \ker H_k \setminus \ker H_j \neq \emptyset$ for all $0 \leq j \leq r$;
- (3) $\bigcap_{k=0, k \neq j}^r (\mathfrak{m}_{\xi_k} \cap R_{\leq d}) \setminus (\mathfrak{m}_{\xi_j} \cap R_{\leq d}) \neq \emptyset$ for all $0 \leq j \leq r$;
- (4) the natural quotient map $R_{\leq d} \rightarrow \bigoplus_{j=0}^r R_{\leq d} / (\mathfrak{m}_{\xi_j} \cap R_{\leq d})$ is surjective;
- (5) $\dim(R_{\leq d} / (\bigcap_{j=0}^r \mathfrak{m}_{\xi_j} \cap R_{\leq d})) = r + 1$;
- (6) $\text{rk}(\sum_{j=0}^r \lambda_j H_j) = r + 1$ for all $\lambda_0, \dots, \lambda_r \in \mathbb{C} \setminus \{0\}$;
- (7) $\text{rk}(\sum_{j=0}^r \lambda_j H_j) = r + 1$ for some $\lambda_0, \dots, \lambda_r \in \mathbb{C} \setminus \{0\}$;
- (8) the pencil (H'_0, \dots, H'_r) is regular and has $r + 1$ different eigenvalues, where $H'_j := (\text{ev}_{\xi_j}(wv))_{w, v \in B}$ for any set of polynomials $B \subseteq R_{\leq d}$ forming a basis of the space $R_{\leq d} / (\bigcap_{j=0}^r \mathfrak{m}_{\xi_j} \cap R_{\leq d})$.

Note that, with respect to the monomial basis of the domain $R_{\leq d}$ and the natural basis of the codomain, the map in (4) corresponds to a multivariate Vandermonde matrix associated to the points ξ_0, \dots, ξ_r .

Proof. The equivalence of (1) and (2) is a special case of Lemma 4.3.1. Moreover, as both $\ker H_k$ and $(\mathfrak{m}_{\xi_k})_{\leq d}$ are of codimension 1 for all $d \in \mathbb{N}$, we have the equality $\ker H_k = (\mathfrak{m}_{\xi_k})_{\leq d}$ for all $0 \leq k \leq r$, which shows equivalence with (3).

Now assume that (3) holds and fix a $j \in \{0, \dots, r\}$, so we can choose a polynomial $p \in R_{\leq d}$ such that $p \in \bigcap_{k=0, k \neq j}^r (\mathfrak{m}_{\xi_k})_{\leq d} \setminus (\mathfrak{m}_{\xi_j})_{\leq d}$. This means that $p(\xi_k) = 0$ for all $k \neq j$ and $p(\xi_j) \neq 0$. Since j was arbitrary, we conclude that there exists a set of Lagrange polynomials in $R_{\leq d}$ that maps to a basis of the $r + 1$ dimensional space $\bigoplus_{j=0}^r R_{\leq d} / (\mathfrak{m}_{\xi_j})_{\leq d}$, so (4) holds. Conversely, let $p \in R_{\leq d}$ be the preimage of the unit vector $e_j \in \bigoplus_{j=0}^r R_{\leq d} / (\mathfrak{m}_{\xi_j})_{\leq d}$ for some $j \in \{0, \dots, r\}$. Then $p(\xi_j) \neq 0$ and $p(\xi_k) = 0$ for all $k \neq j$. Hence, we have $p \notin (\mathfrak{m}_{\xi_j})_{\leq d}$ and $p \in (\mathfrak{m}_{\xi_k})_{\leq d}$ for all $k \neq j$, implying (3).

Next, observe that the map in (4) naturally factors as

$$R_{\leq d} \longrightarrow R_{\leq d} / \left(\bigcap_{j=0}^r \mathfrak{m}_{\xi_j} \cap R_{\leq d} \right) \longrightarrow \bigoplus_{j=0}^r R_{\leq d} / (\mathfrak{m}_{\xi_j} \cap R_{\leq d}).$$

Therefore, as $R / (\bigcap_{j=0}^r \mathfrak{m}_{\xi_j})$ is a finite-dimensional vector space of dimension $r + 1$, it follows that (4) and (5) are equivalent.

For the remaining cases, define the Vandermonde matrix $V = (\xi_j^\alpha)_{0 \leq j \leq r, |\alpha| \leq d}$ such that

$$\sum_{j=0}^r \lambda_j H_j = V^\top \operatorname{diag}(\lambda_0, \dots, \lambda_r) V.$$

As the matrix V corresponds to the map in (4) with respect to the monomial basis, its surjectivity implies that the map $V^\top \operatorname{diag}(\lambda_0, \dots, \lambda_r)$ is injective, so $\operatorname{rk}(\sum_{j=0}^r \lambda_j H_j) = \operatorname{rk} V = r + 1$, which shows that (4) implies (6).

Assume now that (7) holds, which is a special case of (6). In order to show that this implies (8), first observe that we always have the inclusion

$$\ker \left(\sum_{j=0}^r \lambda_j H_j \right) \supseteq \bigcap_{j=0}^r \mathfrak{m}_{\xi_j} \cap R_{\leq d}. \quad (4.9)$$

We claim that, in fact, this is an equality. The space on the right has codimension $\leq r + 1$ with respect to $R_{\leq d}$, so equality follows from the hypothesis that $\ker(\sum_{j=0}^r \lambda_j H_j)$ has codimension $r + 1$, exactly. Therefore, we have $\ker \sum_{j=0}^r \lambda_j H'_j = 0$, so the matrix pencil (H'_0, \dots, H'_r) is regular.

Since $\operatorname{rk}(\sum_{j=0}^r \lambda_j H_j) = r + 1$ by hypothesis, then due to the inclusion (4.9) this means that $\operatorname{rk}(\sum_{j=0}^r \lambda_j H'_j) = r + 1$ as well. As $\ker H'_k$ has codimension 1 for all $0 \leq k \leq r$, it follows that $U_j := \bigcap_{k=0, k \neq j}^r \ker H'_k$ has at most codimension r , so $\dim U_j \geq 1$ for all $0 \leq j \leq r$. Thus, the spaces U_0, \dots, U_r are contained in distinct eigenspaces of the regular pencil (H'_0, \dots, H'_r) , as explained in detail in the proof of Theorem 4.3.5. By Lemma 4.2.9, there cannot exist more than $r + 1$ eigenvalues, as the matrix $\sum_{j=0}^r \lambda_j H'_j$ has rank $r + 1$. Hence, (8) holds.

Finally, assume that (H'_0, \dots, H'_r) is a regular pencil of $m \times m$ -matrices with $r + 1$ different eigenvalues, for some $m \in \mathbb{N}$. As the corresponding eigenvectors must be linearly independent by Lemma 4.2.9, we must have $m \geq r + 1$. However, as m cannot exceed the dimension of $R / (\bigcap_{j=0}^r \mathfrak{m}_{\xi_j})$, which is of dimension $r + 1$ as a vector space, it follows that $m = r + 1$ and therefore (8) implies (5). \square

Remark 4.3.10. The equivalences in Lemma 4.3.9 only hold in a zero-dimensional setting and illustrate that, in the context of positive-dimensional varieties, the conditions (1) to (3) serve as a replacement (for instance, in Algorithm 4.1) for other common descriptions that only hold for the union of finitely many points.

For example, there is no meaningful positive-dimensional interpretation of the Vandermonde matrix being surjective, as in (4), for a variety that consists of infinitely many points. Likewise, the coordinate ring of a positive-dimensional variety is infinite-dimensional as a vector space. The dimension of the truncated coordinate ring, as in (5), is the *affine Hilbert function*. Evaluating it for several degrees can reveal the dimension of the variety (cf. [CLO15, Theorem 9.3.8]), but a property like (5) cannot hold. For a similar reason, a matrix of the form as in (6) usually has larger rank than $r + 1$, under the assumption of conditions (1) to (3) in a positive-dimensional setting.

The description based on eigenvalues as in (8) does not directly transfer to the positive-dimensional case either, since, in that case, the matrix pencil can have more than $r + 1$ eigenvalues. However, in Section 4.4 as well as Section 4.5, we develop criteria that assert that the number of eigenvalues satisfying certain extra conditions does not exceed $r + 1$. \diamond

The following proposition serves as illustration of a direct application of Algorithm 4.1 to the zero-dimensional situation. The statement can be further improved, as is discussed in Remark 4.3.13.

Proposition 4.3.11. *Let $\xi_0, \dots, \xi_r \in \mathbb{C}^n$ be distinct points and $\lambda_0, \dots, \lambda_r \in \mathbb{C} \setminus \{0\}$. Let $H_j = \left(\xi_j^{\alpha+\beta} \right)_{|\alpha|, |\beta| \leq d}$, $0 \leq j \leq r$, for some $d \in \mathbb{N}$ large enough so that the condition $\bigcap_{k=0, k \neq j}^r \ker H_k \setminus \ker H_j \neq \emptyset$ is satisfied. Let $\sigma = \sum_{j=0}^r \lambda_j \text{ev}_{\xi_j}$ and define the shifted matrices*

$$M^{(\nu)} := \left(\sigma(x^{\alpha+\beta+\nu}) \right)_{|\alpha|, |\beta| \leq d} = \sum_{j=0}^r \xi_j^\nu \lambda_j H_j, \quad \nu \in \mathbb{N}^n.$$

Let $\nu_0 := 0, \nu_1, \dots, \nu_s \in \mathbb{N}^n$ for some $s \geq r$ and assume that the matrix $(\xi_j^{\nu_i})_{ij} \in \mathbb{C}^{(s+1) \times (r+1)}$ is of rank $r + 1$. Given $M^{(\nu_0)}, \dots, M^{(\nu_s)}$, then ξ_0, \dots, ξ_r and $\lambda_0, \dots, \lambda_r$ can be recovered using Algorithm 4.1.

Proof. Before being able to apply Algorithm 4.1, we need to construct a regular matrix pencil from the given matrices. For this, note that each matrix H_j has rank 1, so the matrices $M^{(\nu_i)}$ have rank at most $r + 1$. By Lemma 4.3.9 (6), the matrix $M^{(0)} = \sum_{j=0}^r \lambda_j H_j$ has rank exactly $r + 1$.

Thus, we may choose a set of polynomials $B \subseteq R$ that forms a basis of $R_{\leq d} / (\mathfrak{a}_{\leq d} \cap R_{\leq d})$, where $\mathfrak{a} := \bigcap_{j=0}^r \mathfrak{m}_{\xi_j}$, for example by selecting $r + 1$ monomials of degree at most d such that the corresponding $(r + 1) \times (r + 1)$ -submatrix of $M^{(0)}$ still has rank $r + 1$.

More generally, with respect to the basis B , we define the $(r + 1) \times (r + 1)$ -matrices

$$\begin{aligned} H'_j &:= \left(\text{ev}_{\xi_j}(wv) \right)_{w, v \in B}, \quad 0 \leq j \leq r, \\ M'_i &:= \sum_{j=0}^r \xi_j^{\nu_i} \lambda_j H'_j = \left(\sigma(x^{\nu_i} wv) \right)_{w, v \in B}, \quad 0 \leq i \leq s, \end{aligned}$$

where the matrices M'_i can be constructed from the given matrices $M^{(\nu_i)}$. As explained above, the matrix M'_0 has rank $r + 1$, so the pencil (M'_0, \dots, M'_s) is regular. Moreover, the matrix $(\xi_j^{\nu_i})_{ij}$ has full rank by hypothesis, so the same holds for the weighting matrix $(\xi_j^{\nu_i} \lambda_j)_{ij} = (\xi_j^{\nu_i})_{ij} \text{diag}(\lambda_0, \dots, \lambda_r)$.

We are now in a position to apply Algorithm 4.1 to the symmetric matrices M'_0, \dots, M'_s . By Proposition 4.2.20 (2), the pencils (M'_0, \dots, M'_s) and (H'_0, \dots, H'_r) have the same eigenvectors. Hence, due to Lemma 4.3.9 (8), there are exactly $r + 1$ distinct eigenspaces, which implies that the check in Line 8 does not fail.

Thus, we can use Algorithm 4.1 to obtain $[H'_j]$ and $[\xi_j^{\nu_0} \lambda_j : \dots : \xi_j^{\nu_s} \lambda_j]$ for $0 \leq j \leq r$. We can choose representatives \tilde{H}'_j and (c_{0j}, \dots, c_{sj}) , for which we can assume that

$$\sum_{j=0}^r c_{ij} \tilde{H}'_j = \sum_{j=0}^r \xi_j^{\nu_i} \lambda_j H'_j = M'_i, \quad 0 \leq i \leq s,$$

holds. Since we know that $\xi_j^{\nu_0} = 1$, we can further assume that $c_{ij} = \xi_j^{\nu_i}$ and consequently $\tilde{H}'_j = \lambda_j H'_j$. From this, we uniquely recover ξ_j and λ_j for $0 \leq j \leq r$ using the fact that $\ker \tilde{H}'_j = \ker H'_j = (\mathbf{m}_{\xi_j})_{\leq d} / \mathfrak{a}_{\leq d}$. \square

Remark 4.3.12. The assumption in Proposition 4.3.11 that $\bigcap_{k=0, k \neq j}^r \ker H_k \setminus \ker H_j \neq \emptyset$, $0 \leq j \leq r$, is satisfied is equivalent to the property that the Vandermonde matrix $(\xi_j^\alpha)_{0 \leq j \leq r, |\alpha| \leq d}$ is of rank $r + 1$, as follows from Lemma 4.3.9. Thus, if we choose the shifts ν_0, \dots, ν_s suitably, this guarantees that the other condition of Proposition 4.3.11 is fulfilled, namely that the matrix $(\xi_j^{\nu_i})_{0 \leq i \leq s, 0 \leq j \leq r}$ is of rank $r + 1$. This holds when $\{\nu_0, \dots, \nu_s\} \supseteq \mathbb{N}_{\leq d}^n$, in which case one has $r \leq \binom{n+d}{n} - 1 \leq s$. However, this is merely a sufficient condition that is not quite sharp, in many situations, and as such may not represent an optimal choice. A better bound can be obtained by choosing the indices from a hyperbolic cross Υ_r , as in Remark 1.3.8. \diamond

Remark 4.3.13. The conventional multivariate Prony method using matrix pencils requires the matrix $M^{(0)}$ as well as the shifted moment matrices $M^{(e_1)}, \dots, M^{(e_n)}$, where $e_1, \dots, e_n \in \mathbb{N}^n$ denote the standard basis elements. Compared to this, Proposition 4.3.11 needs fewer moments if $s < n$ (for suitable choice of ν_1, \dots, ν_s), but more if $s > n$, which is certainly the case when $r > n$. Moreover, Proposition 4.3.11 is only applicable if the matrix $(\xi_j^{\nu_i})_{ij} \in \mathbb{C}^{(s+1) \times (r+1)}$ is of full rank, which is a further restriction.

Let us more closely investigate the relationship to Prony's method. Due to the structure of the problem, we can refine the process for the case that $s = n < r$ by using only the indices $\nu_0 := 0, \nu_1 := e_1, \dots, \nu_s := e_n \in \mathbb{N}^n$. These are exactly the indices needed for the ordinary matrix pencil method (cf. Section 4.1). This represents an optimal choice of the indices ν_0, \dots, ν_s , as it requires the least amount of moments. Since $s < r$, the requirements of Algorithm 4.1 are not satisfied though, as the matrix $(\xi_j^{\nu_i} \lambda_j)_{ij}$ is of rank smaller than $r + 1$, so the algorithm cannot be applied directly. However, we know that

the matrix $M'_0 = \sum_{j=0}^r \lambda_j H'_j$ is non-singular by construction, so the pencil (M'_0, \dots, M'_s) is regular. Then equation (4.7) shows that $[\xi_j^{\nu_0} \lambda_j : \dots : \xi_j^{\nu_s} \lambda_j] \in \mathbb{P}_{\mathbb{C}}^s$ is an eigenvalue of the pencil (M'_0, \dots, M'_s) for every $0 \leq j \leq r$, as the coordinate $\xi_j^{\nu_0} \lambda_j = \lambda_j$ is non-zero. For this eigenvalue, we can uniquely choose the coordinates $(1, \xi_{j1}, \dots, \xi_{jn}) \in \mathbb{C}^{n+1}$. As the pencil (M'_0, \dots, M'_s) has no more than $r + 1$ eigenvalues, we obtain exactly the points $\xi_0, \dots, \xi_r \in \mathbb{C}^n$ from its eigenvalues. Alternatively, Remark 4.3.8 can be used to accomplish this.

Since $s < r$, we cannot solve the linear system in Line 10 of Algorithm 4.1. However, in this example, the points ξ_0, \dots, ξ_r , that are retrieved as eigenvalues, already completely determine the matrices H_0, \dots, H_r . Thus, we can recover the weights $\lambda_0, \dots, \lambda_r$ by solving a Vandermonde linear system as in Algorithm 1.1. This explains why, in this particular situation, full recovery is possible from fewer shifted matrices than stated in Proposition 4.3.11. \diamond

Example 4.3.14. We now consider the situation of point clusters, in the univariate case $n = 1$. Let $\xi_{jk} \in \mathbb{C}$, $|\xi_{jk}| = 1$, $1 \leq j \leq r$, $1 \leq k \leq m$, be distinct points on the unit circle, where $r \geq 2$, and let $\mu_{jk} \in \mathbb{C} \setminus \{0\}$ be complex weights. Set $d = (r - 1)m$ and define the matrices

$$H_j = \left(\sum_{k=1}^m \mu_{jk} \xi_{jk}^{\alpha+\beta} \right)_{|\alpha|, |\beta| \leq d} \in \mathbb{C}^{(d+1) \times (d+1)}$$

for $1 \leq j \leq r$. Then each H_j is a complex symmetric matrix of rank m .

Next, let $(\lambda_{ij})_{ij} \in \mathbb{C}^{r \times r}$ be an invertible matrix and define

$$M_i = \sum_{j=1}^r \lambda_{ij} H_j$$

for $1 \leq i \leq r$. There exist linear combinations of the matrices H_j , $1 \leq j \leq r$, that are of full rank, which implies that the matrix pencil (M_1, \dots, M_r) is regular, as $(\lambda_{ij})_{ij}$ is invertible.

By choice of d , we have that $U_j := \bigcap_{k=1, k \neq j}^r \ker H_k$ is at least one-dimensional for every $1 \leq j \leq r$. Thus, the requirements for Algorithm 4.1 are satisfied and we may apply it to the matrix pencil (M_1, \dots, M_r) in order to obtain the matrices H_j up to scaling, as long as Line 8 succeeds. Considering $\ker H_j$, this is enough to recover the points $\xi_{j1}, \dots, \xi_{jm}$ for $1 \leq j \leq r$ using the ordinary Prony method. In a final step, the weights μ_{jk} may be recovered from the given data.

Here, we cannot rule out that Line 8 might fail, but for $r \geq 3$ this does not seem to happen in practice with generic input.

A special situation occurs when $r = 2$. In this case, we consider only a pencil (M_1, M_2) of two matrices. This means that the eigenspaces in Line 2 are just the generalized eigenspaces of this pencil, in the usual sense. We know that there are the two eigenvalues $[\lambda_{1j} : \lambda_{2j}]$, $j = 1, 2$, which usually occur with multiplicity 1, but we are unable to

distinguish them from all the other generalized eigenvalues, of which there may be $d + 1$ in total. Hence, Line 8 usually fails.

If in our approach we set $d := (r - 1)m + 1$ instead, it follows that U_j is at least 2-dimensional and, thus, the eigenvalues $[\lambda_{1j} : \lambda_{2j}]$ have multiplicity at least 2 for every $j = 1, 2$. If all the other eigenvalues have smaller multiplicity (or, otherwise, we increase d further), this gives a method to detect the $[\lambda_{1j} : \lambda_{2j}]$ among all the generalized eigenvalues. Then, the subsequent recovery of the points ξ_{jk} works as outlined above. \diamond

Remark 4.3.15. In Example 4.3.14, we may also work with Toeplitz matrices

$$H_j = \left(\sum_{k=1}^m \mu_{jk} \xi_{jk}^{-\alpha+\beta} \right)_{|\alpha|, |\beta| \leq d}$$

if all the points ξ_{jk} lie on the complex torus. These matrices are not Hermitian (unless the μ_{jk} are real), but they satisfy a condition that allows us to adapt Algorithm 4.1 to this situation.

Let $q \in \bigcap_{j=2}^r \ker H_j \setminus \ker H_1$ be an eigenvector of the pencil (H_1, \dots, H_r) . Denote by $V_j := (\xi_{jk}^\beta)_{k\beta}$ the Vandermonde matrices, so that we have $H_j = V_j^* \text{diag}(\mu_{j1}, \dots, \mu_{js}) V_j$ for all j . Then $H_j^* = V_j^* \text{diag}(\overline{\mu_{j1}}, \dots, \overline{\mu_{js}}) V_j$ and it follows that $q \in \ker V_j$ and thus $q \in \ker H_j^*$ for $2 \leq j \leq r$.

Now, if p is another eigenvector of the pencil such that $p \in \bigcap_{j=1, j \neq l}^r \ker H_j$ for some $l \neq 1$ with an eigenvalue $\gamma \in \mathbb{P}_{\mathbb{C}}^{r-1}$, we must have $\gamma_l \neq 0$. It then follows from Corollary 4.2.17 (4) that $q^* H_j p = 0$ for all $1 \leq j \leq r$, so checking for orthogonality as in Line 7 of Algorithm 4.1 is still applicable. \diamond

As another instructive application, we demonstrate how the computation of generalized eigenvalues of a matrix pencil can solve a simple tensor decomposition problem in case of a generic tensor of small rank.

Example 4.3.16. Let $r, s, n, m \in \mathbb{N}$, $1 \leq r \leq \min\{s, n, m\}$, and let \mathbb{k} be an infinite field. Assume that $u_1, \dots, u_r \in \mathbb{k}^s$, $v_1, \dots, v_r \in \mathbb{k}^n$ and $w_1, \dots, w_r \in \mathbb{k}^m$ are generic vectors and define the third-order tensor

$$\sum_{j=1}^r u_j \otimes v_j \otimes w_j \in \mathbb{k}^s \otimes \mathbb{k}^n \otimes \mathbb{k}^m. \quad (4.10)$$

We consider the slices

$$M_i := \sum_{j=1}^r u_{ji} \cdot H_j, \quad 1 \leq i \leq s,$$

of this tensor, where $H_j := v_j \otimes w_j$, $1 \leq j \leq r$, are $n \times m$ -matrices of rank 1. Since the vectors u_j and the rank-1 matrices H_j are generic and \mathbb{k} is infinite, the intersections of the left and right kernels of the matrices M_1, \dots, M_s are of codimension r and we can construct a regular pencil (M'_1, \dots, M'_s) of $r \times r$ -matrices of the form $M'_i = \sum_{j=1}^r u_{ji} H'_j$,

$1 \leq i \leq s$, where the matrices H'_1, \dots, H'_r are $r \times r$ -matrices of rank 1 that form a regular pencil as well (cf. Proposition 4.2.11). In particular, it holds that $\bigcap_{k=1, k \neq j}^r \ker H'_k \neq 0$.

As the vectors u_1, \dots, u_r are linearly independent due to the genericity assumption, we can apply Algorithm 4.1 as explained in Remark 4.3.6. Since the pencil (M'_1, \dots, M'_s) can have at most r eigenvalues by Lemma 4.2.9, the algorithm cannot fail. Thus, we retrieve the eigenvalues $[u_j] = [u_{j1} : \dots : u_{js}] \in \mathbb{P}_{\mathbb{k}}^{s-1}$, $1 \leq j \leq r$, which, up to scaling and permutation, agree with the original vectors u_1, \dots, u_r .

Additionally, we obtain the rank-1 matrices $\tilde{H}'_j \in \mathbb{k}^{r \times r}$ which up to scaling agree with H'_j , $1 \leq j \leq r$. By switching back to the larger vector spaces $\mathbb{k}^n, \mathbb{k}^m$, we then find $H_j \in \mathbb{k}^{n \times m}$ (up to scaling) and recover the remaining vectors $v_j, w_j \in \mathbb{k}^n$ corresponding to u_j , which again are only determined up to scaling.

In summary, computing the eigenvalues from the slices of the tensor (4.10) yields a decomposition of the tensor. The rank of the tensor in this example is r which is less or equal to s, n, m , the dimensions of the vector spaces involved. Though, notice that in general tensors can have larger rank in which case we cannot find a decomposition by computing the eigenvalues like this. We also remark that this approach treats all the slices in an equivalent manner, without the need to single out one or two particular slices that are handled specially. \diamond

Note that, in general, a tensor can have a much higher rank than considered in Example 4.3.16, in which case finding a decomposition is more difficult; see [Lan12, Chapter 12] for an overview. A classic method for this problem is known under the names PARAFAC and CANDECOMP, which is eigenvalue-based as well. It is more general in that it can also be applied to certain tensors of rank r with $2 \leq s < r \leq n, m$. See e.g. [LRA93].

4.3.2 A positive-dimensional counter-example

Here we develop an example of positive-dimensional varieties for which we can show that recovery by Algorithm 4.1 fails if the measures supported on the varieties satisfy certain symmetry conditions.

Lemma 4.3.17. *Let $d \in \mathbb{N}$, let $\text{char } \mathbb{k} \neq 2$ and let $\sigma: \mathbb{k}[x, y]/\langle x^2 + y^2 - 1 \rangle \rightarrow \mathbb{k}$ be a \mathbb{k} -linear map, which is invariant under the ring homomorphisms $(x, y) \mapsto (-x, y)$ and $(x, y) \mapsto (x, -y)$. Then there exists a non-zero univariate polynomial $p(x)$ of degree at most d such that $\sigma(x^\alpha y^\beta p(x)) = 0$ for all $\alpha, \beta \in \mathbb{N}$ with $\alpha + \beta \leq d - 1$.*

Proof. Due to the symmetries, we have $\sigma(x^\alpha y^\beta) = (-1)^\alpha \sigma(x^\alpha y^\beta) = (-1)^\beta \sigma(x^\alpha y^\beta)$ for any $\alpha, \beta \in \mathbb{N}$, which implies $\sigma(x^\alpha y^\beta) = 0$ whenever α or β is odd, as \mathbb{k} is not of characteristic 2. Therefore, it is enough to consider the image of σ on the subalgebra

$$\mathbb{k}[x^2, y^2]/\langle x^2 + y^2 - 1 \rangle = \mathbb{k}[u, v]/\langle u + v - 1 \rangle \cong \mathbb{k}[u],$$

where we identify $u = x^2$ and $v = y^2$. Also note that there is nothing to show if $d = 0$ and that we can choose $p(x) = x$ if $d = 1$, since then $\sigma(p(x)) = \sigma(x) = 0$ due to the symmetries. For the remainder, let us assume that $d \geq 2$.

We construct a polynomial $q(u)$ such that $p(x) = q(x^2)$ if d is even and $xp(x) = q(x^2)$ if d is odd. In the latter case, this means in particular that the constant term of q must be 0. We claim that $\sigma(u^\alpha v^\beta q(u)) = 0$ for all $\alpha, \beta \in \mathbb{N}$ with $\alpha + \beta \leq \lfloor \frac{d}{2} \rfloor - 1$. In particular, if d is even, this implies that $\sigma(x^{2\alpha} y^{2\beta} p(x)) = 0$ for $2\alpha + 2\beta \leq d - 1$, while for odd d it implies $\sigma(x^{2\alpha} y^{2\beta} xp(x)) = 0$ for $2\alpha + 2\beta + 1 \leq d - 1$. Hence, the statement follows from the claim.

To prove the claim, it is enough to show $\sigma(u^\alpha q(u)) = 0$ for $0 \leq \alpha \leq \lfloor \frac{d}{2} \rfloor - 1$, since $v \equiv 1 - u \pmod{\langle u + v - 1 \rangle}$. As $q(u)$ is a polynomial of degree at most $\lfloor \frac{d}{2} \rfloor$, we can construct q from the non-trivial kernel of the Hankel matrix

$$\left(\sigma(u^{\alpha+\alpha'}) \right)_{0 \leq \alpha \leq \lfloor \frac{d}{2} \rfloor - 1, 0 \leq \alpha' \leq \lfloor \frac{d}{2} \rfloor}.$$

Due to the dimensions of this matrix, its kernel is at least one-dimensional if d is even and two-dimensional if d is odd. This means that, for odd d , we can choose q in such a way that the additional linear requirement of its constant term being 0 is satisfied. \square

Proposition 4.3.18. *Let d, \mathbb{k}, σ and p be as in Lemma 4.3.17 and define*

$$H_j = \left(\sigma((x, y - \xi_j)^{\alpha+\beta}) \right)_{|\alpha|, |\beta| \leq d}$$

for distinct $\xi_j \in \mathbb{k}$, $0 \leq j \leq r$. Then $H_0 p = \dots = H_r p$.

In particular, p is an eigenvector of the pencil (H_0, \dots, H_r) if the pencil is regular.

Proof. In order to show $H_0 p = \dots = H_r p$, it is enough to assert that

$$\sigma((x, y)^\alpha p(x, y)) = \sigma((x, y + \xi)^\alpha p(x, y + \xi))$$

holds for all $|\alpha| \leq d$ and arbitrary translations $\xi \in \mathbb{k}$. As, by construction, p is constant in y , this is equivalent to showing

$$\sigma(x^{\alpha_1} ((y + \xi)^{\alpha_2} - y^{\alpha_2}) p(x)) = 0$$

for $|\alpha| \leq d$. For $\alpha_2 = 0$, this is trivial. Otherwise, we have

$$\sigma(x^{\alpha_1} (y + \xi^{\alpha_2} - y^{\alpha_2}) p(x)) = \sum_{k=0}^{\alpha_2-1} \binom{\alpha_2}{k} \xi^{\alpha_2-k} \sigma(x^{\alpha_1} y^k p(x)),$$

which vanishes by Lemma 4.3.17.

The addendum follows from Proposition 4.2.5 (3), as $p \neq 0$ by Lemma 4.3.17. \square

Example 4.3.19. We construct an example corresponding to positive-dimensional varieties for which Algorithm 4.1 fails. For this, we consider $r + 1$ different circles in \mathbb{R}^2 of equal radius such that the centers of the circles lie on a line. Without loss of generality, we can assume that the radius is 1 and that the line is a coordinate axis. We consider the moments with respect to the uniform measure on these circles. For this, let $R := \mathbb{C}[x, y]$ and let $\sigma: R \rightarrow \mathbb{C}$ be the complexification of the moment functional of the unit circle. Note that σ satisfies the invariance properties of Lemma 4.3.17.

Denote by $\mathfrak{a}_j := \langle x^2 + (y - \xi_j)^2 - 1 \rangle$, $0 \leq j \leq r$, the vanishing ideals of the circles. Assume that $d \in \mathbb{N}$ such that $2r \leq d < 2(r + 1)$. Then for every $0 \leq j \leq r$, we find a non-zero polynomial of degree at most d in $\bigcap_{k=0, k \neq j}^r \mathfrak{a}_k \neq 0$ which is not contained in $\mathfrak{a}_j \cap R_{\leq d}$.

The corresponding moment matrices H_j , $0 \leq j \leq r$, are of the form as in Proposition 4.3.18. It follows from Theorem 3.4.11 that H_j induces a positive-definite matrix on the truncated coordinate ring $R_{\leq d}/(\mathfrak{a}_j \cap R_{\leq d})$, which means that $\ker H_j = \mathfrak{a}_j \cap R_{\leq d}$, for all $0 \leq j \leq r$. Thus, we have

$$\bigcap_{\substack{k=0 \\ k \neq j}}^r \ker H_k = \bigcap_{\substack{k=0 \\ k \neq j}}^r \mathfrak{a}_k \cap R_{\leq d} \neq 0.$$

Additionally, by choice of d , we have

$$\bigcap_{k=0}^r \ker H_k = \bigcap_{k=0}^r \mathfrak{a}_k \cap R_{\leq d} = 0,$$

so that the pencil (H_0, \dots, H_r) is regular by Lemma 4.2.13. This means the conditions on H_j for Algorithm 4.1 are satisfied.

Thus, with a matrix $(\lambda_{ij})_{ij} \in \mathbb{C}^{(s+1) \times (r+1)}$ of rank $r + 1$, where $s \geq r$, and $M_i := \sum_{j=0}^r \lambda_{ij} H_j$, $0 \leq i \leq s$, we may attempt to apply Algorithm 4.1. We obtain the $r + 1$ distinct points $[\lambda_{0j} : \dots : \lambda_{sj}] \in \mathbb{P}_{\mathbb{C}}^s$, $0 \leq j \leq r$, as eigenvalues of the pencil (M_0, \dots, M_s) . In addition, by Proposition 4.3.18, the polynomial p constructed in Lemma 4.3.17 occurs as an eigenvector with eigenvalue $[\sum_{k=0}^r \lambda_{0k} : \dots : \sum_{k=0}^r \lambda_{sk}] \in \mathbb{P}_{\mathbb{C}}^s$, as it is the image of the eigenvalue $[1 : \dots : 1] \in \mathbb{P}_{\mathbb{C}}^r$ of the pencil (H_0, \dots, H_r) under the map induced by $(\lambda_{ij})_{ij}$, by Proposition 4.2.20. Since $(\lambda_{ij})_{ij}$ has full rank, we have

$$[\lambda_{0j} : \dots : \lambda_{sj}] \neq [\sum_{k=0}^r \lambda_{0k} : \dots : \sum_{k=0}^r \lambda_{sk}]$$

for all $0 \leq j \leq r$. Therefore, we always obtain at least $r + 2$ different spaces in Line 8 of Algorithm 4.1, so that the recovery fails.

If the centers of the circles do not lie on a line, Proposition 4.3.18 does not apply and, in practice, one indeed only obtains $r + 1$ different spaces in Line 8 for generic input, so that the recovery algorithm works. \diamond

Example 4.3.19 illustrates that, as a consequence of Proposition 4.2.20, we can think of the cases in which recovery by Algorithm 4.1 succeeds as exactly those cases in which the matrices H_j , $0 \leq j \leq r$, have the property that the eigenvalues of the pencil (H_0, \dots, H_r) are exactly the $r + 1$ unit vectors. In other words, the eigenvectors are contained in the kernel of all but one matrix H_j . Note, in particular, that this does not depend on the weights λ_{ij} , but only on the matrices H_j . We will come back to this example with a more explicit computation later, in Example 4.8.12.

4.4 Positive-semidefinite case

In the case of positive-semidefinite matrices, it is possible to infer additional information about the eigenvalues of a matrix pencil. We will show that the subset of eigenvalues we want to reconstruct are the vertices of the convex polytope that is spanned by the set of all eigenvalues, as long as we make a suitable choice of coordinates. Based on that, we can formulate a more advanced algorithm that resolves the shortcomings of Algorithm 4.1. As we work over the complex numbers here, the convex polytopes we consider live in \mathbb{C}^{r+1} and are the image of real convex polytopes under a \mathbb{C} -linear map.

Before presenting the new algorithm, we first take note of a few general results for pencils of positive-semidefinite matrices. Let us start with the following consequence of Lemma 4.2.13, pertaining to the choice of coordinates for the eigenvalues.

Corollary 4.4.1. *Let (H_0, \dots, H_r) be a regular pencil of positive-semidefinite matrices in $\mathbb{C}^{n \times n}$. If $\gamma \in \mathbb{P}_{\mathbb{C}}^r$ is an eigenvalue of the pencil, then one can choose real non-negative coordinates for γ .*

Proof. By the proof of Lemma 4.2.13, the matrix $H := \sum_{j=0}^r \lambda_j H_j$ is non-singular for arbitrary $\lambda_0, \dots, \lambda_r > 0$. Assume that $p \in \mathbb{C}^n$, $p \neq 0$, is an eigenvector of the pencil for the eigenvalue γ . Then we can choose coordinates $\gamma_0, \dots, \gamma_r$ for γ such that $H_j p = \gamma_j H p$ for all $0 \leq j \leq r$. Thus, we have $p^* H_j p \geq 0$, as H_j is positive-semidefinite, as well as $p^* H p > 0$, since H is non-singular, so we conclude that $\gamma_j \geq 0$, for all $0 \leq j \leq r$. \square

Note, however, that a regular pencil of Hermitian matrices may have non-real eigenvalues, despite the fact that the ordinary eigenvalues of an individual Hermitian matrix are real; see for instance Example 4.2.19.

Remark 4.4.2. Corollary 4.4.1 shows that we can choose non-negative coordinates for the eigenvalues of a regular pencil of positive-semidefinite matrices. In particular, we can normalize them such that they sum to 1.

On the other hand, we can uniquely choose coordinates with respect to a fixed non-singular matrix H which is a linear combination of the matrices H_0, \dots, H_r . If we choose $H = \sum_{j=0}^r H_j$, which is non-singular by the proof of Lemma 4.2.13, then by

Corollary 4.2.7 the respective coordinates of an eigenvalue $\gamma \in \mathbb{P}_{\mathbb{C}}^r$ are

$$\frac{1}{\sum_{j=0}^r \gamma_j}(\gamma_0, \dots, \gamma_r),$$

so they sum to 1, as well. A direct consequence of this is that, in these coordinates, the eigenvalue is a convex combination of the $r+1$ unit vectors. More generally, the following property holds for a concrete choice of coordinates of an eigenvalue. \diamond

Lemma 4.4.3. *Let H_0, \dots, H_r be positive-semidefinite matrices in $\mathbb{C}^{n \times n}$ and let $H = \sum_{j=0}^r c_j H_j$ be any linear combination with $c_0, \dots, c_r \in \mathbb{R}_{>0}$. Let $\gamma \in \mathbb{C}^{r+1}$ and $p \in \mathbb{C}^n \setminus \ker H$ such that $H_j p = \gamma_j H p$ for all $0 \leq j \leq r$. Then*

$$\gamma \in \text{conv} \left\{ \frac{e_0}{c_0}, \dots, \frac{e_r}{c_r} \right\},$$

where $e_0, \dots, e_r \in \mathbb{C}^{r+1}$ denote the standard basis vectors.

Note that $[\gamma] \in \mathbb{P}_{\mathbb{C}}^r$ is an eigenvalue if the pencil (H_0, \dots, H_r) is regular. Under the additional assumption that the matrices H_0, \dots, H_r satisfy $\bigcap_{k=0, k \neq j}^r \ker H_k \setminus \ker H_j \neq \emptyset$ for all $0 \leq j \leq r$, then also the points $[e_0], \dots, [e_r] \in \mathbb{P}_{\mathbb{C}}^r$ are eigenvalues of the pencil, as the elements in $\bigcap_{k=0, k \neq j}^r \ker H_k \setminus \ker H_j$, $0 \leq j \leq r$, are corresponding eigenvectors. With the above choice of coordinates, these eigenvalues are represented by $\frac{e_0}{c_0}, \dots, \frac{e_r}{c_r}$ in \mathbb{C}^{r+1} . Hence, in this case, every eigenvalue of the pencil is a convex combination of the representing coordinates of this distinguished set of $r+1$ eigenvalues.

Proof. First observe that H is a positive-semidefinite matrix, as it is a linear combination of positive-semidefinite matrices with positive coefficients c_0, \dots, c_r . As $p \notin \ker H$, we therefore have $p^* H p > 0$ and it follows that

$$\gamma_j = \frac{p^* H_j p}{p^* H p} \geq 0,$$

for all $0 \leq j \leq r$. By definition of H , this also implies that

$$\sum_{j=0}^r c_j \gamma_j = \sum_{j=0}^r \frac{c_j p^* H_j p}{p^* H p} = 1.$$

Moreover, we have

$$\gamma = \sum_{j=0}^r \gamma_j e_j = \sum_{j=0}^r c_j \gamma_j \frac{e_j}{c_j}.$$

Since $c_j \gamma_j \geq 0$, this is a convex combination of $\frac{e_0}{c_0}, \dots, \frac{e_r}{c_r}$. \square

By a similar proof, we obtain the following result, which extends Lemma 4.4.3 to the case of linear transformations.

4.4 POSITIVE-SEMIDEFINITE CASE

Lemma 4.4.4. *Let H_0, \dots, H_r be positive-semidefinite matrices and let $H = \sum_{j=0}^r c_j H_j$ be any linear combination with $c_0, \dots, c_r \in \mathbb{R}_{>0}$. Let $s \geq 0$ and let $M_i = \sum_{j=0}^r \lambda_{ij} H_j$ for $0 \leq i \leq s$, where $(\lambda_{ij})_{ij} \in \mathbb{C}^{(s+1) \times (r+1)}$ is an arbitrary matrix. Further, define $\eta_j := \frac{1}{c_j}(\lambda_{0j}, \dots, \lambda_{sj}) \in \mathbb{C}^{s+1}$ for $0 \leq j \leq r$. Then the following properties hold.*

(1) *If $\gamma \in \mathbb{C}^{s+1}$ such that $M_i p = \gamma_i H p$, $0 \leq i \leq s$, for some vector $p \notin \ker H$, then*

$$\gamma \in \text{conv}\{\eta_0, \dots, \eta_r\}.$$

(2) *If the matrix $(\lambda_{ij})_{ij}$ is of rank $r + 1$, then, for every $0 \leq k \leq r$, the point η_k is a non-redundant point of the convex hull, so*

$$\eta_k \notin \text{conv}\{\eta_0, \dots, \eta_{k-1}, \eta_{k+1}, \dots, \eta_r\}.$$

(3) *If $\eta_0 \notin \text{conv}\{\eta_1, \dots, \eta_r\}$ and if q is any vector, the following are equivalent:*

(a) $q \in \bigcap_{j=1}^r \ker H_j \setminus \ker H_0$,

(b) $q \notin \ker H$ and $M_i q = \eta_{0i} H q$ for all $0 \leq i \leq s$.

Proof. As H is positive-semidefinite and $H p \neq 0$, we have $p^* H p > 0$, so it follows for $0 \leq i \leq s$ that

$$\gamma_i = \frac{p^* M_i p}{p^* H p} = \sum_{j=0}^r \frac{\lambda_{ij}}{c_j} a_j, \quad (4.11)$$

where we define $a_j := c_j \frac{p^* H_j p}{p^* H p}$ for $0 \leq j \leq r$. These parameters satisfy $a_j \geq 0$ and $a_0 + \dots + a_r = 1$, so, in view of (4.11), we infer that the equation

$$\gamma = \sum_{j=0}^r a_j \eta_j$$

expresses γ as a convex combination of η_0, \dots, η_r , which proves (1).

For (2), assume that $\eta_0 \in \text{conv}\{\eta_1, \dots, \eta_r\}$, which means that

$$\text{conv}\{\eta_1, \dots, \eta_r\} = \text{conv}\{\eta_0, \dots, \eta_r\}. \quad (4.12)$$

However, $\text{conv}\{\eta_0, \dots, \eta_r\}$ is the image of the convex polytope $\text{conv}\{\frac{e_0}{c_0}, \dots, \frac{e_r}{c_r}\}$ under the map defined by $(\lambda_{ij})_{ij}$; cf. Lemma 4.4.3. By hypothesis, this map is injective, so the image of the polytope is of real dimension r and cannot be defined by less than $r + 1$ points, contradicting (4.12). Thus, we conclude that $\eta_0 \notin \text{conv}\{\eta_1, \dots, \eta_r\}$ and the result follows.

For (3), let us assume that $q \in \bigcap_{j=1}^r \ker H_j \setminus \ker H_0$. Then in particular $q \notin \ker H$, since $c_0 \neq 0$. Additionally, we have for all $0 \leq i \leq s$ that

$$M_i q = \lambda_{i0} H_0 q = \eta_{0i} c_0 H_0 q = \eta_{0i} H q,$$

by definition of η_0 .

For the converse, we first show that $q \in \ker H_j$ for $1 \leq j \leq r$. By the previous part of the proof, we can write η_0 as a convex combination

$$\eta_0 = \sum_{j=0}^r a_j \eta_j,$$

with parameters $a_j := c_j \frac{q^* H_j q}{q^* H q} \geq 0$, $0 \leq j \leq r$, for which $a_0 + \dots + a_r = 1$. Assuming that $a_0 \neq 1$, it follows that $\eta_0 = \frac{1}{1-a_0} \sum_{j=1}^r a_j \eta_j$. As $\sum_{j=1}^r \frac{a_j}{1-a_0} = 1$, this is a contradiction to the hypothesis that $\eta_0 \notin \text{conv}\{\eta_1, \dots, \eta_r\}$. Therefore, we must have $a_0 = 1$ and $a_1 = \dots = a_r = 0$, which implies that $q \in \ker H_1, \dots, \ker H_r$. Finally, observe that assuming $q \in \ker H_0$ leads to a contradiction to the hypothesis that $q \notin \ker H$, so we conclude that $q \notin \ker H_0$. \square

Remark 4.4.5. Note that the statements (1) and (3) of Lemma 4.4.4 do not assume that the matrix $(\lambda_{ij})_{ij}$ is of rank $r+1$, unlike results with similar hypotheses in other sections. In particular, the case $s=0$ is allowed, in which we consider only a single Matrix M_0 and η_0, \dots, η_r are just complex numbers. Moreover, it is not necessary that H is invertible or that (H_0, \dots, H_r) is a regular matrix pencil, which would be a requirement for the notion of eigenvalues we work with. \diamond

Now we can refine our previous reconstruction algorithm, Algorithm 4.1, for the case of positive-semidefinite matrices. This is specified in Algorithm 4.2.

Algorithm 4.2 Positive-semidefinite recovery from several weighted sums

Input: A natural number $s \in \mathbb{N}$ and a regular pencil (M_0, \dots, M_s) of complex matrices.

Assumptions: The matrices are of the form $M_i = \sum_{j=0}^r \lambda_{ij} H_j$, $0 \leq i \leq s$, for some $r \leq s$, where H_0, \dots, H_r are positive-semidefinite matrices and where $(\lambda_{ij})_{ij} \in \mathbb{C}^{(s+1) \times (r+1)}$ is a matrix of rank $r+1$. Additionally, we have that $\lambda_{0j} \in \mathbb{R}_{>0}$ and $\bigcap_{k=0, k \neq j}^r \ker H_k \neq 0$, for every $0 \leq j \leq r$.

Output: r as well as $[\lambda_{0j} : \dots : \lambda_{sj}]$ for all $0 \leq j \leq r$ (up to scaling and permutation).

- 1: Let $H := M_0$, which is a positive-definite matrix.
 - 2: Compute the set of eigenvalues of the matrix pencil (M_0, \dots, M_s) . Denote by $\Gamma \subseteq \mathbb{C}^{s+1}$ the set of their representatives, by choosing coordinates with respect to the matrix H (cf. Corollary 4.2.7).
 - 3: Compute the set of vertices of the convex polytope $\text{conv } \Gamma$ and denote it by Γ' .
 - 4: Set $r := \#\Gamma' - 1$.
 - 5: With $\Gamma' = \{\gamma_0, \dots, \gamma_r\}$, we have $[\lambda_{0j} : \dots : \lambda_{sj}] = [\gamma_j]$ for $0 \leq j \leq r$ (up to permutations in j).
-

Remark 4.4.6. In contrast to Algorithm 4.1, note that Algorithm 4.2 cannot fail. Any eigenvalues that would lead to the failure of Algorithm 4.1 are filtered out in this refined algorithm by using the additional piece of information that those particular eigenvalues cannot be extreme points of the convex hull of all eigenvalues of the pencil.

It is important that the coordinates for the eigenvalues in Line 2 are chosen with respect to the single fixed matrix H . The algorithm crucially depends on the fact that H is positive-definite. In fact, if we were to choose new coordinates with respect to a different non-singular matrix, the new coordinates would, in general, not be expressible as the product of a fixed scalar and the original coordinates, so the computation of the eigenvalues needs to make sure to choose the coordinates correctly and consistently.

If we assume that $\lambda_{ij} \geq 0$ for all i, j , we can drop the assumption that $\lambda_{0j} > 0$ for all $0 \leq j \leq r$. In this case, all the M_i are positive-semidefinite, so we can choose H as $H = \sum_{i=0}^s M_i$ instead, for example. Indeed, this matrix is positive-semidefinite and it is regular by Lemma 4.2.13, since the pencil (M_0, \dots, M_s) is regular.

More precisely, observe that the assumption $\lambda_{0j} > 0$ for $0 \leq j \leq r$ is equivalent to assuming that M_0 is positive-definite. This follows from the hypothesis that $\bigcap_{k=0, k \neq j}^r \ker H_k \neq 0$ for all $0 \leq j \leq r$. As such, it is enough to replace H by any positive-definite linear combination of M_0, \dots, M_s . However, without further assumptions on the coefficients λ_{ij} , this can be difficult to find in general. The tools of semidefinite programming may provide a possible approach for this problem. See [VB96, Section 6] for an overview of the real situation.

If the pencil (M_0, \dots, M_s) is singular, then, since the matrices are linear combinations of positive-semidefinite matrices and since the matrix $(\lambda_{ij})_{ij}$ has full column rank, it follows from Lemma 4.2.13 that one can construct a regular pencil from this by dividing out the common kernel of the matrices. This preprocessing step therefore allows to apply Algorithm 4.2 in the case of a singular pencil as well. This is explained in more detail in the proof of Theorem 4.7.4, as Algorithm 4.3 makes use of this.

In a postprocessing step, we may reconstruct the matrices H_0, \dots, H_r (up to scaling) or the kernels of these matrices, by the computation in Line 10 of Algorithm 4.1. \diamond

Theorem 4.4.7. *Under the given assumptions, recovery by Algorithm 4.2 works.*

Proof. The proof of Lemma 4.2.13 shows that any linear combination of H_0, \dots, H_r with positive coefficients is a non-singular matrix. Therefore, $H = M_0 = \sum_{j=0}^r \lambda_{0j} H_j$ is positive-definite, as claimed.

Define $\gamma_j := \frac{1}{\lambda_{0j}}(\lambda_{0j}, \dots, \lambda_{sj}) \in \mathbb{C}^{s+1}$. In particular, we have the equality $[\gamma_j] = [\lambda_{0j} : \dots : \lambda_{sj}]$ in $\mathbb{P}_{\mathbb{C}}^s$, for all $0 \leq j \leq r$, as is claimed in Line 5. Then, by Lemma 4.4.4 (1), it follows for the set Γ computed in Line 2, which represents the eigenvalues of the pencil (M_0, \dots, M_s) in terms of coordinates with respect to the matrix H , that

$$\Gamma \subseteq \text{conv}\{\gamma_0, \dots, \gamma_r\}.$$

Denote by $e_0, \dots, e_r \in \mathbb{C}^{r+1}$ the standard basis vectors. By Lemma 4.3.1, we have $\bigcap_{k=0, k \neq j}^r \ker H_k \setminus \ker H_j \neq \emptyset$ for every $0 \leq j \leq r$, since the kernels of H_0, \dots, H_r intersect trivially, as the pencil (H_0, \dots, H_r) is regular. Therefore, $[e_0], \dots, [e_r] \in \mathbb{P}_{\mathbb{C}}^r$ are eigenvalues of the pencil (H_0, \dots, H_r) and, as a consequence of Proposition 4.2.20,

$[\gamma_0], \dots, [\gamma_r] \in \mathbb{P}_{\mathbb{C}}^s$ are eigenvalues of (M_0, \dots, M_s) . For the representing coordinates, this implies by Remark 4.2.21 that

$$\gamma_0, \dots, \gamma_r \in \Gamma,$$

as the coordinates are unique by Corollary 4.2.7, and thus $\text{conv } \Gamma = \text{conv}\{\gamma_0, \dots, \gamma_r\}$. As the points $\gamma_0, \dots, \gamma_r$ are non-redundant by Lemma 4.4.4(2), they must be the vertices of the polytope $\text{conv } \Gamma$. Hence, we have $\Gamma' = \{\gamma_0, \dots, \gamma_r\}$. \square

4.5 A sufficient condition

In this section, we develop a sufficient criterion for Algorithm 4.1, or variants thereof, to succeed. The new algorithms based on this criterion are introduced in Section 4.7.

We now focus on the situation of moment problems on algebraic varieties, for which we use the framework from Section 3.1 in order to handle both the affine and torus case. Hence, in the following, we assume that the field \mathbb{k} and the rings R, L are chosen as in Section 3.1. The following remark introduces the main concepts of this section in terms of a two-component example.

Remark 4.5.1. Consider two \mathbb{k} -linear maps $\sigma_1, \sigma_2: L \rightarrow \mathbb{k}$ with ideals $\mathfrak{a}_j \subseteq L$ such that $\mathfrak{a}_j, \mathfrak{a}_j^\circ \subseteq \ker \sigma_j$ for $j = 1, 2$. Let $B \subseteq R$ be polynomials that form a basis of the quotient space $R_{\leq d}/(\mathfrak{a}_1 \cap \mathfrak{a}_2 \cap R_{\leq d})$. Denote by H_1, H_2 the Gramian matrices describing the sesquilinear forms on the quotient space with respect to the basis B which are induced by the sesquilinear forms associated to σ_1, σ_2 as explained by Lemma 3.1.11, i. e. we have $H_j = (\sigma_j(w^\circ v))_{w, v \in B}$, $j = 1, 2$. Then the matrix pencil (H_1, H_2) is regular if and only if the induced sesquilinear form of some linear combination of σ_1, σ_2 is non-degenerate (or equivalently, right-non-degenerate) on $R_{\leq d}/(\mathfrak{a}_1 \cap \mathfrak{a}_2 \cap R_{\leq d})$.

This means that, for generic choices of $\gamma \in \mathbb{k}$, the matrix $H_1 - \gamma H_2$ is regular. This is equivalent to saying that, if there exists a polynomial $p \in R_{\leq d}$ satisfying

$$\sigma_1(q^\circ p) = \gamma \sigma_2(q^\circ p) \tag{4.13}$$

for all $q \in R_{\leq d}$, then we already have $p \in \mathfrak{a}_1 \cap \mathfrak{a}_2$, by Lemma 3.1.17.

However, for a particular (non-generic) choice of γ , this need not hold – namely for the eigenvalues of the pencil (H_1, H_2) . In that case, the corresponding p satisfying (4.13) is an eigenvector of the pencil, unless of course p is zero. Indeed, this becomes apparent in terms of the basis B , so write p as $p = \sum_{w \in B} p_w w$ with coefficients $p_w \in \mathbb{k}$. Then Equation (4.13) gives

$$H_1 p = (\sigma_1(w^\circ v))_{w, v \in B} (p_v)_{v \in B} = \gamma (\sigma_2(w^\circ v))_{w, v \in B} (p_v)_{v \in B} = \gamma H_2 p,$$

showing that p is an eigenvector. In particular, the vector space

$$(\mathfrak{a}_1 \cap R_{\leq d})/(\mathfrak{a}_1 \cap \mathfrak{a}_2 \cap R_{\leq d})$$

is a subspace of the eigenspace with $\gamma = 0$, since, if $p \in \mathfrak{a}_1$, then also $q^\circ p \in \mathfrak{a}_1 \subseteq \ker \sigma_1$ for every q . Therefore, its vector space dimension is a lower bound for the dimension of the corresponding eigenspace. In many cases we consider, these dimensions are in fact equal. \diamond

As the size of the matrices in Remark 4.5.1 is $\dim(R_{\leq d}/(\mathfrak{a}_1 \cap \mathfrak{a}_2 \cap R_{\leq d}))$, this implies that the maximal number of extraneous eigenvalues, i. e. eigenvalues that do not correspond to eigenvectors in the spaces

$$(\mathfrak{a}_j \cap R_{\leq d})/(\mathfrak{a}_1 \cap \mathfrak{a}_2 \cap R_{\leq d}), \quad j = 1, 2, \quad (4.14)$$

is bounded by the quantity in Lemma 1.2.4 (2), if we assume that the hypotheses of the lemma are satisfied. Hence, if $s \in \mathbb{N}$ is the dimension of the components, then, by Lemma 1.2.4, the dimension of the spaces (4.14) is in $\Theta(d^s)$, while the maximum number of extraneous eigenvalues is in $O(d^{s-1})$, which is of strictly smaller order. Similarly, this holds for more than two components, as well. Note that this also explains why in the zero-dimensional case, where $s = 0$, such a problem does not arise.

As we are primarily interested in finding criteria that rule out the existence of extraneous eigenvalues, let us expand on Remark 4.5.1. The discussion motivates the following lemma, which, under slightly stronger assumptions, gives further insight into the nature of the eigenvectors of a pencil as in Remark 4.5.1. It yields a sufficient condition for asserting that extraneous eigenvectors as in Algorithm 4.1 and illustrated in Example 4.3.19 do not exist, which is made precise in Theorem 4.5.4. The stronger hypothesis states that not only p , but also certain polynomial multiples hp occur as eigenvectors for the eigenvalue γ of the pencil.

Lemma 4.5.2. *Let $\sigma_1, \sigma_2: L \rightarrow \mathbb{k}$ be \mathbb{k} -linear maps and let $\mathfrak{a}_1, \mathfrak{a}_2 \subseteq L$ be ideals such that $\mathfrak{a}_j \subseteq \ker \sigma_j$, $j = 1, 2$. Let $d, \delta \in \mathbb{N}$ and assume that there exists a linear combination σ of σ_1, σ_2 such that the induced sesquilinear map $\Phi_{\mathfrak{a}_1 \cap \mathfrak{a}_2}^{R_{\leq d+\delta}}$ (defined in Definition 3.1.12) associated to σ is right-non-degenerate.*

Let $\gamma \in \mathbb{k}$ and $p \in R_{\leq d}$ such that

$$\sigma_1(g^\circ hp) = \gamma \sigma_2(g^\circ hp)$$

for all $g \in R_{\leq d+\delta}$ and $h \in R_{\leq \delta}$. Additionally, assume that one of the following conditions holds for the ideals $\mathfrak{a}_1, \mathfrak{a}_2$:

- (1) \mathfrak{a}_1 is prime and there exists a polynomial $q \in \mathfrak{a}_2 \cap R_{\leq \delta} \setminus \mathfrak{a}_1$;
- (2) $\mathfrak{a}_1 = \mathfrak{p}_1 \cap \cdots \cap \mathfrak{p}_m$ for prime ideals $\mathfrak{p}_1, \dots, \mathfrak{p}_m \subseteq L$ and there exists a polynomial $q \in \mathfrak{a}_2 \cap R_{\leq \delta} \setminus \bigcup_{i=1}^m \mathfrak{p}_i$;
- (3) \mathfrak{a}_1 is prime with $\mathfrak{a}_1 \not\supseteq \mathfrak{a}_2$ and \mathfrak{a}_2 is generated by $\mathfrak{a}_2 \cap R_{\leq \delta}$;
- (4) $\mathfrak{a}_1 = \mathfrak{p}_1 \cap \cdots \cap \mathfrak{p}_m$ for prime ideals $\mathfrak{p}_i \subseteq L$ with $\mathfrak{p}_i \not\supseteq \mathfrak{a}_2$, $1 \leq i \leq m$, and the ideals \mathfrak{a}_2 and \mathfrak{p}_i are generated by $\mathfrak{a}_2 \cap R_{\leq d_0}$ and $\mathfrak{p}_i \cap R_{\leq d_i}$, respectively, for $d_0, d_i \in \mathbb{N}$, $1 \leq i \leq m-1$, such that $\sum_{i=0}^{m-1} d_i \leq \delta$;

(5) $\mathfrak{a}_1, \mathfrak{a}_2$ are comaximal and \mathfrak{a}_2 is generated by $\mathfrak{a}_2 \cap R_{\leq \delta}$.

Then $p \in \mathfrak{a}_1$.

Proof. Case (1) is a special case of (2), so we assume that (2) is satisfied. Then it follows from $q \in \mathfrak{a}_2 \subseteq \ker \sigma_2$ and $q \in R_{\leq \delta}$ that

$$\sigma_1(g^\circ qp) = \gamma \sigma_2(g^\circ qp) = 0$$

for all $g \in R_{\leq d+\delta}$. In particular, this means that $\sigma(g^\circ qp) = 0$ for all $g \in R_{\leq d+\delta}$, so Lemma 3.1.17 implies that $qp \in \mathfrak{a}_1 \cap \mathfrak{a}_2 \cap R_{\leq d+\delta}$. In particular, we have $qp \in \mathfrak{p}_i$ for every $1 \leq i \leq m$. Since \mathfrak{p}_i is prime and $q \notin \mathfrak{p}_i$, we conclude that $p \in \mathfrak{p}_i$ for every $1 \leq i \leq m$ and therefore $p \in \mathfrak{a}_1$.

Case (3) is a special case of (4) with $m = 1$.

For (4), we use prime avoidance to construct an element $q \in \mathfrak{a}_2 \setminus \bigcup_{i=1}^m \mathfrak{p}_i$ with $q \in R_{\leq \delta}$, so that the statement follows from (2). Note that we can assume that no inclusions exist among the prime ideals, so, for all $1 \leq i, j \leq m$, we have $\mathfrak{p}_i \not\subseteq \mathfrak{p}_j$, since otherwise \mathfrak{p}_j would be redundant in the decomposition of \mathfrak{a}_1 as an intersection of prime ideals.

We use induction over m to construct the polynomial q . The case $m = 0$ is trivial. If $m = 1$, it follows from $\mathfrak{a}_2 \not\subseteq \mathfrak{p}_1$ that some generator q of \mathfrak{a}_2 is not contained in \mathfrak{p}_1 . As \mathfrak{a}_2 is generated by elements in $R_{\leq d_0}$, we can assume that $q \in R_{\leq d_0} \subseteq R_{\leq \delta}$.

Next, assume that $m > 1$ and that the inductive hypothesis is satisfied for $m - 1$. Thus, since in particular $\sum_{i=0}^{m-2} d_i \leq \sum_{i=0}^{m-1} d_i \leq \delta$, we find a polynomial $q \in \mathfrak{a}_2 \setminus \bigcup_{i=1}^{m-1} \mathfrak{p}_i$ with $q \in R_{\leq \delta}$. If $q \notin \mathfrak{p}_m$, then q already has the desired property, so we assume that $q \in \mathfrak{p}_m$. As \mathfrak{p}_m is a prime ideal and since $\mathfrak{a}_2, \mathfrak{p}_1, \dots, \mathfrak{p}_{m-1} \not\subseteq \mathfrak{p}_m$ by assumption, it follows by Lemma 1.2.3 that

$$\mathfrak{a}_2 \mathfrak{p}_1 \cdots \mathfrak{p}_{m-1} \not\subseteq \mathfrak{p}_m.$$

Since $\sum_{i=0}^{m-1} d_i \leq \delta$, the ideal product $\mathfrak{a}_2 \mathfrak{p}_1 \cdots \mathfrak{p}_{m-1}$ is generated by elements in $R_{\leq \delta}$ and, additionally, not all generators are contained in \mathfrak{p}_m . Hence, we can choose an element $q' \in \mathfrak{a}_2 \mathfrak{p}_1 \cdots \mathfrak{p}_{m-1} \setminus \mathfrak{p}_m$ such that $q' \in R_{\leq \delta}$. Then $q + q'$ is a polynomial in $R_{\leq \delta}$ that is also contained in $\mathfrak{a}_2 \setminus \bigcup_{i=1}^m \mathfrak{p}_i$.

In case (5), fix an arbitrary element $q \in \mathfrak{a}_2 \cap R_{\leq \delta}$. Similarly to case (2), it follows that $\sigma(g^\circ qp) = 0$ for all $g \in R_{\leq d+\delta}$ and therefore $qp \in \mathfrak{a}_1 \cap \mathfrak{a}_2$ by Lemma 3.1.17. As \mathfrak{a}_2 is generated by $\mathfrak{a}_2 \cap R_{\leq \delta}$ and since q was an arbitrary element from $\mathfrak{a}_2 \cap R_{\leq \delta}$, this implies that $\langle p \rangle \mathfrak{a}_2 \subseteq \mathfrak{a}_1 \cap \mathfrak{a}_2$. Hence, as $\mathfrak{a}_1, \mathfrak{a}_2$ are comaximal, this means that $p \in \mathfrak{a}_1$ by Lemma 1.2.1. \square

Remark 4.5.3. The proof of Lemma 4.5.2 shows that we can formulate cases (1) and (2) under the weaker assumption that $\sigma_1(g^\circ hp) = \gamma \sigma_2(g^\circ hp)$ for $g \in R_{\leq d+\delta}$ and for $h = q$. In practice, this property is difficult to check without knowing q explicitly, which is why the formulation of the lemma seems more useful, in most situations. \diamond

Theorem 4.5.4. *Let $\sigma_j: L \rightarrow \mathbb{k}$, $0 \leq j \leq r$, be \mathbb{k} -linear maps and let $\mathfrak{a}_j \subseteq L$ be ideals such that $\mathfrak{a}_j \subseteq \ker \sigma_j$. Let $d, \delta \in \mathbb{N}$ and assume that, for every $0 \leq i < j \leq r$, the induced sesquilinear map $\Phi_{\mathfrak{a}_i \cap \mathfrak{a}_j}^{R_{\leq d+\delta}}$ associated to some linear combination of σ_i, σ_j is right-non-degenerate.*

Let $\gamma = [\gamma_0 : \cdots : \gamma_r] \in \mathbb{P}_{\mathbb{k}}^r$ and $p \in R_{\leq d}$ such that, for all $0 \leq i, j \leq r$,

$$\gamma_i \sigma_j(g^\circ hp) = \gamma_j \sigma_i(g^\circ hp) \quad (4.15)$$

holds for all $g \in R_{\leq d+\delta}, h \in R_{\leq \delta}$ and such that, for each pair of ideals $\mathfrak{a}_i, \mathfrak{a}_j$ with $i \neq j$, one of the conditions (1)–(5) listed in Lemma 4.5.2 is satisfied.

If $p \notin \mathfrak{a}_0 \cap \cdots \cap \mathfrak{a}_r$, then there exists an index $k \in \{0, \dots, r\}$ such that $p \in \bigcap_{j=0, j \neq k}^r \mathfrak{a}_j$ and $\gamma_j = 0$ for all $0 \leq j \leq r$ with $j \neq k$.

Proof. Since $\gamma \in \mathbb{P}_{\mathbb{k}}^r$, we can choose k such that $\gamma_k \neq 0$. Without loss of generality, we can assume that $k = 0$. Then

$$\sigma_j(g^\circ hp) = \frac{\gamma_j}{\gamma_0} \sigma_0(g^\circ hp)$$

holds for all $g \in R_{\leq d+\delta}, h \in R_{\leq \delta}$ and $j \neq 0$. Thus, by Lemma 4.5.2, it follows that $p \in \mathfrak{a}_j$ for all $j \neq 0$.

It remains to show that $\gamma_j = 0$ for $j \neq 0$. Assuming that $\gamma_j \neq 0$ for some $j \neq 0$, we have $\sigma_0(g^\circ hp) = \frac{\gamma_0}{\gamma_j} \sigma_j(g^\circ hp)$ for $g \in R_{\leq d+\delta}, h \in R_{\leq \delta}$ and may apply Lemma 4.5.2 again in order to deduce that $p \in \mathfrak{a}_0$. This however is a contradiction to the assumption that $p \notin \mathfrak{a}_0 \cap \cdots \cap \mathfrak{a}_r$, so we conclude that $\gamma = [1 : 0 : \cdots : 0]$. \square

Remark 4.5.5. Our main application of Theorem 4.5.4 is the case in which each pair of ideals satisfies property (3) of Lemma 4.5.2: the ideals $\mathfrak{a}_0, \dots, \mathfrak{a}_r$ are prime ideals that are generated by polynomials of degree at most δ and which satisfy $\mathfrak{a}_i \not\subseteq \mathfrak{a}_j$ for all $0 \leq i, j \leq r$ with $i \neq j$. When working with a degree-induced filtration, under this assumption, δ can often be chosen relatively small in comparison to d , since d is a bound on the degree of the polynomial $p \in \bigcap_{j=0, j \neq k}^r \mathfrak{a}_j$ for some $k \in \{0, \dots, r\}$, a polynomial vanishing on many distinct varieties, and therefore commonly needs to be quite large.

Note that condition (3) of Lemma 4.5.2 is a special case of the strictly weaker condition (1). Indeed, every generator of \mathfrak{a}_2 is contained in $R_{\leq \delta}$ and at least one of them must not be contained in \mathfrak{a}_1 in order that $\mathfrak{a}_1 \not\supseteq \mathfrak{a}_2$ can be satisfied. For convenience and simplicity of exposition, we primarily work with condition (3), rather than (1), but it is worth noting that the latter may allow to infer better bounds on δ , in certain examples.

Furthermore, note that the assumption in Theorem 4.5.4 that, for every pair i, j with $0 \leq i < j \leq r$, there exists a linear combination of σ_i, σ_j such that the induced sesquilinear map $\Phi_{\mathfrak{a}_i \cap \mathfrak{a}_j}^{R_{\leq d+\delta}}$ is right-non-degenerate does not necessarily imply that there exists a linear combination σ of $\sigma_0, \dots, \sigma_r$ such that the induced sesquilinear map $\Phi_{\mathfrak{a}_0 \cap \cdots \cap \mathfrak{a}_r, \sigma}^{R_{\leq d+\delta}}$ is right-non-degenerate.

What is more, if we assume for simplicity that the sesquilinear forms $\langle -, - \rangle_{\sigma_j}$ are Hermitian, then we cannot conclude that the pencil (H_0, \dots, H_r) is regular, where H_j denotes the matrix of $\langle -, - \rangle_{\sigma_j}$ with respect to some basis of $R_{\leq d+\delta}/(\mathfrak{a}_0 \cap \dots \cap \mathfrak{a}_r \cap R_{\leq d+\delta})$. Indeed, a generic linear combination $\sum_{j=0}^r c_j H_j$ can be of rank strictly smaller than the dimension of the latter space, as the kernel of such a linear combination can vary with the parameters $c_j \in \mathbb{k}$, $0 \leq j \leq r$, as noted in Example 4.2.12.

However, if $\mathbb{k} \subseteq \mathbb{C}$ and the sesquilinear forms $\langle -, - \rangle_{\sigma_j}$ are positive-semidefinite on $R_{\leq d}$, it does follow that the pencil is regular, as stated in Corollary 4.5.6 below. This is one of the main motivations for focusing on the situation in which these forms are positive-semidefinite, as it allows us to infer the regularity of the pencil. Another motivation is the fact that non-negative measures give rise to positive-semidefinite forms which therefore pose an interesting and relevant class of examples.

Also note that, under the assumption that the pencil (H_0, \dots, H_r) is regular, then, up to a change of basis, the properties of (4.15) in particular describe what it means for p (as well as its polynomial multiples hp for $h \in R_{\leq \delta}$) to be an eigenvector with eigenvalue γ of the pencil; cf. Proposition 4.2.5. \diamond

Corollary 4.5.6. *Assume that $\mathbb{k} \subseteq \mathbb{C}$ with complex conjugation as involution $-^\circ$. Let $d, \delta \in \mathbb{N}$, let $\sigma_j: L \rightarrow \mathbb{k}$, $0 \leq j \leq r$, be \mathbb{k} -linear maps and let $\mathfrak{p}_j \subseteq L$ be prime ideals such that $\mathfrak{p}_j \subseteq \ker \sigma_j$ and $\mathfrak{p}_i \not\subseteq \mathfrak{p}_j$ for all $i \neq j$. Assume that, for each $0 \leq j \leq r$, the ideal \mathfrak{p}_j is generated by $\mathfrak{p}_j \cap R_{\leq \delta}$ and the sesquilinear form $\langle -, - \rangle_{\sigma_j}$ associated to σ_j on $R_{\leq d+\delta}$ is positive-semidefinite.*

Additionally, assume that one of the following conditions holds:

- (1) *the induced sesquilinear form on $R_{\leq d+\delta}/(\mathfrak{p}_j \cap R_{\leq d+\delta})$ associated to σ_j is non-degenerate for every $0 \leq j \leq r$;*
- (2) *the induced sesquilinear form on $R_{\leq d+\delta}/(\mathfrak{p}_i \cap \mathfrak{p}_j \cap R_{\leq d+\delta})$ associated to some linear combination of σ_i, σ_j is non-degenerate for all $0 \leq i < j \leq r$.*

For $0 \leq j \leq r$, denote by H_j the Gramian matrix of the induced sesquilinear form associated to σ_j with respect to some basis of $U := R_{\leq d+\delta}/(\mathfrak{p}_0 \cap \dots \cap \mathfrak{p}_r \cap R_{\leq d+\delta})$. Then the matrix pencil (H_0, \dots, H_r) is regular.

If there exists $\gamma \in \mathbb{P}_{\mathbb{k}}^r$ and an element $\bar{p} \in U$, $\bar{p} \neq 0$, represented by a polynomial $p \in R_{\leq d}$ such that $\bar{q}\bar{p} \in U$ is an eigenvector with eigenvalue γ of the pencil (H_0, \dots, H_r) for all $q \in R_{\leq \delta}$ that satisfy $\bar{q}\bar{p} \neq 0$ in U , then there exists a $k \in \{0, \dots, r\}$ such that $\gamma_j = 0$ for all $0 \leq j \leq r$ with $j \neq k$ and it holds that $p \in \bigcap_{j=0, j \neq k}^r \mathfrak{p}_j$.

In the statement above and in the proof below, \bar{g} denotes the residue class modulo $\mathfrak{p}_0 \cap \dots \cap \mathfrak{p}_r \cap R_{\leq d+\delta}$ of a polynomial $g \in R_{\leq d+\delta}$ and not complex conjugation. Furthermore, note that the sesquilinear forms associated to $\sigma_0, \dots, \sigma_r$ on $R_{\leq d+\delta}$ are positive-semidefinite, so in particular they are Hermitian, which means that the induced sesquilinear forms on the quotient spaces above are well-defined by Lemma 3.1.11. Additionally,

recall that, by Theorem 3.4.11, assumption (1) holds for any moment functional σ_j of a non-negative measure if the ideal \mathfrak{p}_j is the vanishing ideal of its support.

Proof. First, observe that (1) implies (2). Indeed, for every $0 \leq i < j \leq r$, we may consider the matrices H'_i, H'_j describing the induced sesquilinear form of σ_i, σ_j on $R_{\leq d+\delta}/(\mathfrak{p}_i \cap \mathfrak{p}_j \cap R_{\leq d+\delta})$. As these matrices are positive-semidefinite, it follows from Lemma 4.2.13 that (H'_i, H'_j) forms a regular pencil if the kernels of these matrices intersect trivially. Condition (1) implies by Lemmas 3.1.16 and 3.1.17 that the kernels of H'_i, H'_j are represented by elements in $\mathfrak{p}_i, \mathfrak{p}_j$, respectively. Thus, the intersection of $\ker H'_i$ and $\ker H'_j$ is represented by elements in $\mathfrak{p}_i \cap \mathfrak{p}_j$, so, in the space $R_{\leq d+\delta}/(\mathfrak{p}_i \cap \mathfrak{p}_j \cap R_{\leq d+\delta})$, the intersection of the kernels is trivial. This means that there exists a linear combination of H'_i, H'_j which is regular and therefore describes a non-degenerate sesquilinear form. Hence, we may assume that condition (2) is satisfied.

In this case, it follows by a similar argument that the matrix pencil (H_0, \dots, H_r) is regular. Indeed, if there exists a $u \in \bigcap_{j=0}^r \ker H_j$, then u has a representative in $\mathfrak{p}_i \cap \mathfrak{p}_j$ for all $0 \leq i < j \leq r$ by Lemmas 3.1.16 and 3.1.17 and therefore $u = 0$ in U . Thus, the pencil is regular by Lemma 4.2.13.

As the ideals $\mathfrak{p}_0, \dots, \mathfrak{p}_r$ are generated by elements in $R_{\leq \delta}$, each pair of them satisfies condition (3) of Lemma 4.5.2. Then the conclusion follows from Theorem 4.5.4, together with Lemma 3.1.16, if we can show that the properties of (4.15) are satisfied.

Note that $\bar{q}p = 0$ if and only if $\bar{q} = 0$ in U , as $\mathfrak{p}_0, \dots, \mathfrak{p}_r$ are prime ideals and $\bar{p} \neq 0$ by assumption. If $\bar{q}p \neq 0$ is an eigenvector with eigenvalue γ for all $q \in R_{\leq \delta}$ with $\bar{q} \neq 0$, we therefore have

$$\gamma_i \sigma_j(g^\circ qp) = \gamma_j \sigma_i(g^\circ qp)$$

for all $g \in R_{\leq d+\delta}$ and all $q \in R_{\leq \delta}$, as $\mathfrak{p}_j \subseteq \ker \sigma_j$ for all j . This means we can apply Theorem 4.5.4 with degrees d and δ . \square

Example 4.5.7. For illustration, let us apply Theorem 4.5.4 to the zero-dimensional case of distinct points $\xi_0, \dots, \xi_r \in \mathbb{k} := \mathbb{C}$. For simplicity, we focus on the affine situation, so let $L = R = \mathbb{C}[x_1, \dots, x_n]$ with trivial involution and with the filtration induced by total degree. The ideals are of the form $\mathfrak{a}_j := \mathfrak{m}_{\xi_j} = \langle x - \xi_j \rangle$, $0 \leq j \leq r$. For every $0 \leq j \leq r$, the maps $\sigma_j: R \rightarrow \mathbb{C}$ are arbitrary \mathbb{C} -linear maps such that $\mathfrak{a}_j \subseteq \ker \sigma_j$. As discussed in Example 3.2.1, this implies that σ_j is of the form $\sigma_j = \lambda_j \text{ev}_{\xi_j}$ for some $\lambda_j \in \mathbb{C}$, that is, $\sigma_j(p) = \lambda_j p(\xi_j)$ for all $p \in R$. We further assume that σ_j is not the zero-map, which means that $\lambda_j \neq 0$ for all $0 \leq j \leq r$.

The ideals \mathfrak{a}_j are generated by polynomials of degree 1, so we can pick $\delta = 1$. Therefore, as the ideals are pairwise comaximal, condition (5) of Lemma 4.5.2 is satisfied for each pair of the ideals.

Next, let us check the right-non-degenerateness requirement. If $d \in \mathbb{N}$ and $0 \leq i < j \leq r$, we need to show that there exists a linear combination $\sigma := c_i \sigma_i + c_j \sigma_j$ with $c_i, c_j \in \mathbb{C}$ for which the sesquilinear map $\Phi_{\mathfrak{a}_i \cap \mathfrak{a}_j, \sigma}^{R_{\leq d+\delta}}$ is right-non-degenerate. As the involutions are

trivial in this example, this map is a bilinear form on the space $R_{\leq d+\delta}/(\mathfrak{a}_i \cap \mathfrak{a}_j \cap R_{\leq d+\delta})$, which is a two-dimensional vector space as ξ_i and ξ_j differ in at least one coordinate and $R/(\mathfrak{a}_i \cap \mathfrak{a}_j)$ is two-dimensional. This bilinear form is induced by the bilinear form $\langle -, - \rangle_\sigma$ on $R_{\leq d+\delta}$, which we can represent more concretely in terms of the Hankel matrix

$$\left(c_i \lambda_i \xi_i^{\alpha+\beta} + c_j \lambda_j \xi_j^{\alpha+\beta} \right)_{|\alpha|, |\beta| \leq d+\delta}.$$

Clearly, for any $c_i, c_j \in \mathbb{C} \setminus \{0\}$ and $d \in \mathbb{N}$, this matrix has rank 2. Then it follows from the characterization in Lemma 4.3.9 (8) that $\Phi_{\mathfrak{a}_i \cap \mathfrak{a}_j, \sigma}^{R_{\leq d+\delta}}$ is non-degenerate for some choice of c_i, c_j . (Alternatively, this follows from Theorem 3.4.11 applied to the measure $\delta_{\xi_i} + \delta_{\xi_j}$, which has the moment functional $\lambda_i^{-1} \sigma_i + \lambda_j^{-1} \sigma_j$, giving rise to a positive-definite bilinear form on $R_{\leq d+\delta}/(\mathfrak{a}_i \cap \mathfrak{a}_j \cap R_{\leq d+\delta})$.)

Thus, for point ideals $\mathfrak{a}_j = \mathfrak{m}_{\xi_j} = \langle x - \xi_j \rangle$, $0 \leq j \leq r$, all the requirements of Theorem 4.5.4 are satisfied as long as $\lambda_0, \dots, \lambda_r \neq 0$.

Now, if $p \in R_{\leq d}$ satisfies the eigenvalue-eigenvector-type condition (4.15) for some $\gamma \in \mathbb{P}_{\mathbb{C}}^r$, then the statement of Theorem 4.5.4 is that p either vanishes at all the points ξ_0, \dots, ξ_r or it vanishes at all but one of the points. In the latter case, p is a Lagrange-like polynomial in the sense of Remark 4.3.3. Beyond that, no other polynomial can satisfy this condition.

Since non-zero polynomials vanishing at r or $r+1$ points cannot usually exist if the degree d is very small (depending on the point configuration), by contraposition, we conclude that there only exist polynomials satisfying the eigenvector-type condition (4.15) if the degree d is sufficiently large. \diamond

4.6 Sharpness results

So far, it is not clear whether the bounds on the degrees in Lemma 4.5.2, Theorem 4.5.4 or Corollary 4.5.6 are sharp. To this end, we would like to investigate whether, with the notation of Lemma 4.5.2, there exist examples such that $\sigma_1(g^\circ h p) = \gamma \sigma_2(g^\circ h p)$ is satisfied for all $g \in R_{\leq d+\delta-1}$ and $h \in R_{\leq \delta-1}$, while $p \notin \mathfrak{a}_1$ holds. The following provides a result in this direction, for the univariate case where $n = 1$.

Example 4.6.1. Let $\mathbb{k} = \mathbb{C}$ with complex conjugation as involution. Let $R = \mathbb{C}[x]$ and $L = \mathbb{C}[x^{\pm 1}]$ be the univariate polynomial and Laurent polynomial ring. In addition, we define the involution on L as in Example 3.1.5 and choose the degree-induced filtration on R .

Furthermore, let $r \geq 2$ be a natural number and denote by $\xi_j := \exp\left(\frac{2\pi i j}{r}\right) \in \mathbb{T} \subseteq \mathbb{C}$ for $0 \leq j < r$ the r -th roots of unity. Thus, the set of points is described by the ideal

$$\mathfrak{a}_1 := \bigcap_{j=0}^{r-1} \langle x - \xi_j \rangle \subseteq L.$$

4.6 SHARPNESS RESULTS

Additionally, choose $\omega \in \mathbb{C}$ with $|\omega| = 1$ such that ω is not an r -th root of unity and define the ideal

$$\mathfrak{a}_2 := \bigcap_{j=0}^{r-1} \langle x - \omega \xi_j \rangle \subseteq L.$$

The corresponding zero-dimensional varieties are displayed in Figure 4.1. By choice of the parameter ω , the ideals $\mathfrak{a}_1, \mathfrak{a}_2$ are comaximal. Observe that $\mathfrak{a}_1, \mathfrak{a}_2$ are generated by polynomials of degree r , so they are generated by $\mathfrak{a}_1 \cap R_{\leq r}$ and $\mathfrak{a}_2 \cap R_{\leq r}$, respectively.




Figure 4.1: The point sets $V(\mathfrak{a}_1), V(\mathfrak{a}_2)$, denoted by \bullet and \circ , on the unit circle \mathbb{T} , for $r = 3$ and arbitrary parameter ω .

We define the moment functionals

$$\sigma_1: L \longrightarrow \mathbb{C}, \quad x^\alpha \longmapsto \sum_{j=0}^{r-1} \xi_j^\alpha, \quad \sigma_2: L \longrightarrow \mathbb{C}, \quad x^\alpha \longmapsto \sum_{j=0}^{r-1} (\omega \xi_j)^\alpha,$$

for all $\alpha \in \mathbb{Z}$. First, observe that the $r \times r$ -Toeplitz matrices (cf. Remark 3.1.7)

$$T_k := \sigma_k(x^{-\alpha+\beta})_{0 \leq \alpha, \beta \leq r-1}, \quad k = 1, 2,$$

are positive-definite, which can be seen from the Vandermonde decompositions of these matrices. Therefore, the matrix pencil (T_1, T_2) is regular and the induced sesquilinear form on $R_{\leq \delta} / (\mathfrak{a}_1 \cap \mathfrak{a}_2 \cap R_{\leq \delta})$ associated to convex combinations of σ_1, σ_2 is positive-definite, so in particular non-degenerate, for all $\delta \in \{0, \dots, r-1\}$. As the form is Hermitian, also the induced sesquilinear maps $\Phi_{\mathfrak{a}_1 \cap \mathfrak{a}_2}^{R_{\leq \delta}}$ are non-degenerate by Lemmas 3.1.9 and 3.1.16.

Next, let $d = 0$ and let $p = 1 \in R_{\leq d} \subseteq R$, so p is a polynomial of degree 0 which is not contained in the ideal \mathfrak{a}_1 , in particular. We claim that, for any $\delta \in \{0, \dots, r-1\}$, we have

$$\sigma_1(g^\circ hp) = \sigma_2(g^\circ hp)$$

for all $g, h \in R_{\leq \delta}$. Note that, if this were true for $\delta = r$, then, since $d = 0$, the statement of Lemma 4.5.2 (5) would imply that p is contained in \mathfrak{a}_1 , which is a contradiction. In other words, this example shows that the bounds on the degrees in Lemma 4.5.2, Theorem 4.5.4 and Corollary 4.5.6 are sharp, in this case.

To prove the claim, it is enough to show it for the monomial basis of $R_{\leq \delta}$, so we have to prove that

$$\sigma_1(x^{-\alpha+\beta}) = \sigma_2(x^{-\alpha+\beta})$$

for $0 \leq \alpha, \beta \leq \delta$. For this, we need to show that $\sigma_1(x^\alpha) = \sigma_2(x^\alpha)$ for $-\delta \leq \alpha \leq \delta$ or, equivalently, that

$$(\sigma_1 - \sigma_2)(x^\alpha) = \sum_{j=0}^{r-1} \xi_j^\alpha - (\omega \xi_j)^\alpha = (1 - \omega^\alpha) \sum_{j=0}^{r-1} \xi_j^\alpha \quad (4.16)$$

is zero. The factor $(1 - \omega^\alpha)$ vanishes for $\alpha = 0$, the other factor on the right for $\alpha \not\equiv 0 \pmod{r}$ as the points ξ_0, \dots, ξ_{r-1} are the r -th roots of unity. Indeed, this follows from

$$\sum_{j=0}^{r-1} \xi_j^\alpha = \sum_{j=0}^{r-1} \xi_1^{\alpha j} = \frac{\xi_1^{\alpha r} - 1}{\xi_1^\alpha - 1} = 0$$

for $\alpha \not\equiv 0 \pmod{r}$. Hence, since $\delta < r$, the expression (4.16) vanishes in particular for $-\delta \leq \alpha \leq \delta$ as claimed. \diamond

Next, let us elaborate on the results proved in Section 4.5, from a different point of view. We have formulated all the statements in terms of an involution $-\circ$ on L , for which we primarily think of one of the two cases where either the involution extends complex conjugation to the Laurent ring L and R is the polynomial ring as in Example 3.1.5 or the involution is trivial and $L = R$ is the polynomial ring as in Example 3.1.4. In the first case, it is useful to think of Toeplitz moment matrices, whereas one can think of Hankel matrices in the second case, as explained in Remark 3.1.7.

These two cases are not entirely unrelated. As established by Section 3.3, in particular Lemma 3.3.1, a Toeplitz matrix conveys essentially the same information as a Hankel matrix of corresponding size. The converse is not necessarily true – geometrically speaking, this stems from the fact that the natural embedding of the algebraic torus in affine space is not invertible, so that a Toeplitz matrix cannot encode certain information coming from outside the torus.

A natural question that arises from this is whether the results for the Toeplitz case proved in Section 4.5 already follow from the results for the Hankel case. Here, we explain why this is not the case. Thus, the additional complexity introduced by considering these cases simultaneously actually pays off, as either case provides useful new information.

For simplicity, we again focus on the univariate situation where $n = 1$.

Example 4.6.2. We consider two cases.

- (1) Let $R = \mathbb{C}[z]$, $L = \mathbb{C}[z^{\pm 1}]$, and
- (2) let $R' = \mathbb{C}[x, y]$, $S = R'/\langle x^2 + y^2 - 1 \rangle$.

These cases are related via the isomorphism $\varphi: L \xrightarrow{\sim} S$ defined by (3.2) in Section 3.3. Moreover, we fix the filtrations on R and R' that are induced by total degree and fix the involution $-\circ$ on L as defined in Example 3.1.5 as well as a trivial involution on R' .

Note that this means that the involution on R' is not compatible with the involution on L under φ , since the one on L extends complex conjugation of the ground field \mathbb{C} ,

whereas the one on R' does not. This is in contrast to the discussion in Remark 3.3.2, but this distinction is not important for the purpose of this example. Our choice here has the advantage that it allows us to define bilinear forms on the quotient spaces $R'_{\leq d+\delta}/(\mathfrak{q}_i \cap \mathfrak{q}_j \cap R'_{\leq d+\delta})$, as we do below.

We first discuss case (1). Let $r \in \mathbb{N}$ and let $\xi_j \in \mathbb{C}$ with $|\xi_j| = 1$, $0 \leq j \leq r$, be distinct points on the torus \mathbb{T} . We define the maximal ideals $\mathfrak{p}_j := \langle z - \xi_j \rangle \subseteq L$ and the functionals $\sigma_j: L \rightarrow \mathbb{C}$, $z^\alpha \mapsto \xi_j^\alpha$, which are the evaluation homomorphisms at the points ξ_j and to which we associate sesquilinear forms on L . In particular, we have $\mathfrak{p}_j \subseteq \ker \sigma_j$. Additionally, note that \mathfrak{p}_j is generated by $\mathfrak{p}_j \cap R_{\leq 1}$, so we set $\delta := 1$.

We wish to study this example in view of Theorem 4.5.4. One readily checks that \mathfrak{p}_j and σ_j , $0 \leq j \leq r$, satisfy the criteria of that theorem. Now, let $d \in \mathbb{N}$ and assume that $p \in R_{\leq d}$, $p \notin \mathfrak{p}_0 \cap \cdots \cap \mathfrak{p}_r$, is a polynomial satisfying the assumptions of Theorem 4.5.4, for some $\gamma \in \mathbb{P}_{\mathbb{C}}^r$. Then it follows that $p \in \mathfrak{p}_1 \cap \cdots \cap \mathfrak{p}_r$ and $p \notin \mathfrak{p}_0$, where, without loss of generality, we assume that $k = 0$.

The main question we are interested in is what the minimal degree of a polynomial p with such properties can be, so we wish to choose d as small as possible. Clearly, p is divisible by $\prod_{j=1}^r (z - \xi_j)$, so we conclude that $d \geq r$. As the latter polynomial vanishes on all the components defined by $\mathfrak{p}_1, \dots, \mathfrak{p}_r$, there exists a polynomial of degree equal to r with those properties, so we can assume that $d = r$, in the best case.

For case (2) on the other hand, we may consider this example in terms of S or its covering R' . The ideals $\varphi(\mathfrak{p}_j) \subseteq S$ are of the form $\mathfrak{q}_j := \langle x - \Re(\xi_j), y - \Im(\xi_j) \rangle$, but in light of Theorem 4.5.4 it is convenient to view these ideals \mathfrak{q}_j as prime ideals in R' . Denote by $\tau_j: R' \rightarrow \mathbb{C}$, $0 \leq j \leq r$, the \mathbb{C} -linear maps defined by $\tau_j(q) = \sigma_j(\varphi^{-1}(\bar{q}))$ for $q \in R'$, so in particular $\mathfrak{q}_j \subseteq \ker \tau_j$. Again observe that \mathfrak{q}_j is generated by $\mathfrak{q}_j \cap R'_{\leq \delta}$, where $\delta = 1$ as before. In this case, we associate bilinear forms on R' to τ_j , $0 \leq j \leq r$.

Let $d' \in \mathbb{N}$. Then $R'_{\leq d'+\delta}/(\mathfrak{q}_i \cap \mathfrak{q}_j \cap R'_{\leq d'+\delta})$ is a two-dimensional vector space for any $0 \leq i, j \leq r$ with $i \neq j$ and one checks that linear combinations of τ_i, τ_j give rise to non-degenerate bilinear forms on this space; see for instance Lemma 4.3.9 (8).

Now, if $p' \in R'_{\leq d'}$, $p' \notin \mathfrak{q}_0 \cap \cdots \cap \mathfrak{q}_r$, is a polynomial that satisfies the assumptions of Theorem 4.5.4 with respect to \mathfrak{q}_j and τ_j for $0 \leq j \leq r$ and for some $\gamma' \in \mathbb{P}_{\mathbb{C}}^r$, then it follows that $p' \in \mathfrak{q}_1 \cap \cdots \cap \mathfrak{q}_r$ and $p' \notin \mathfrak{q}_0$, assuming that $k = 0$ as above. We are again interested in a lower bound on the degree d' of such a polynomial p' .

As $p' \notin \mathfrak{q}_0 \cap \cdots \cap \mathfrak{q}_r$, it is evident that p' does not vanish identically on the circle. However, it vanishes at the r points $(\Re(\xi_j), \Im(\xi_j))$, $1 \leq j \leq r$, which lie on the circle. As a consequence of Bézout's theorem, it follows that $2 \deg(p') \geq r$, so we conclude that $d' \geq \lceil \frac{r}{2} \rceil$. Since any two points on the circle are uniquely cut out by a line, it is clear that there actually exists a polynomial in R' of degree equal to $\lceil \frac{r}{2} \rceil$ that vanishes exactly at those r points, so we can assume that $d' = \lceil \frac{r}{2} \rceil$, in the best case.

Finally, let us compare the two cases. Working with the monomial bases as in Remark 3.1.7, we can assume that the relevant data in case (2) is given in terms of the

Hankel matrices

$$(\tau_j((x, y)^{\alpha+\beta}))_{\alpha, \beta \in \mathbb{N}^2, |\alpha|, |\beta| \leq d'+\delta},$$

for $0 \leq j \leq r$. By Lemma 3.3.1, these matrices convey the same information as the Toeplitz matrices

$$T'_j = (\sigma_j(z^{-k+l}))_{-(d'+\delta) \leq k, l \leq d'+\delta} = (\sigma_j(z^{-k+l}))_{0 \leq k, l \leq 2d'+2\delta}$$

(see also Remark 3.1.14). Similarly, in case (1), we can assume that the data is given in terms of the Toeplitz matrices

$$T_j = (\sigma_j(z^{-k+l}))_{0 \leq k, l \leq d+\delta}.$$

From here, we see that the matrices T_j are smaller than T'_j . More precisely, the difference between the degree bounds is

$$(2d' + 2\delta) - (d + \delta) = 2\lceil \frac{r}{2} \rceil - r + \delta = \begin{cases} 1 & \text{if } r \text{ is even,} \\ 2 & \text{if } r \text{ is odd.} \end{cases}$$

Thus, the formulation in case (1) allows for a sharper and thus stronger result than in case (2). \diamond

Example 4.6.2 can be transferred to the multivariate setting where $n > 1$ by embedding \mathbb{T}^1 in \mathbb{T}^n . This means that there exist certain point configurations in \mathbb{T}^n , e. g. when all points lie in a one-dimensional subspace, for which the degree of the polynomial p satisfying the properties of Theorem 4.5.4 (as well as the size of the involved moment matrices) can be chosen to be smaller in the Toeplitz formulation than in the corresponding Hankel formulation.

4.7 Revised algorithms

The sufficient condition developed in Section 4.5 allows us to make refinements to Algorithm 4.1, for the case of moment matrices of measures supported on algebraic varieties. Based on that, we present two new algorithms here, Algorithms 4.3 and 4.4, that allow for successful recovery of the underlying components.

In this section, we assume that \mathbb{k} is \mathbb{R} or \mathbb{C} . On the one hand, this has practical reasons because algorithms derived here rely on the numerical computation of eigenvalues, a task most commonly performed over the real or complex numbers. On the other hand, we assume that the sesquilinear forms associated to the \mathbb{k} -linear maps σ_j in Algorithm 4.3 are positive-semidefinite, which limits the reasonable choices of \mathbb{k} ; see Definition 3.1.6. Recall that one of the advantages of the sesquilinear forms being positive-semidefinite is that they can be viewed as sesquilinear forms on the quotient spaces $R_{\leq d}/(\mathfrak{p}_j \cap R_{\leq d})$ as long as $\mathfrak{p}_j \subseteq \ker \sigma_j$ by Lemma 3.1.11. Therefore, in this section, we require that one of the following two hypotheses holds:

- (1) $\mathbb{k} = \mathbb{R}$, $L = R = \mathbb{R}[x_1, \dots, x_n]$ with trivial involutions (cf. Example 3.1.4), or
- (2) $\mathbb{k} = \mathbb{C}$, $R = \mathbb{C}[x_1, \dots, x_n]$, $L = \mathbb{C}[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ with complex conjugation and involution $-\circ$ on L defined as in Example 3.1.5.

For the filtrations on R , we choose one that is induced by total degree or max-degree. This conveniently allows us to pick the monomials as bases for the vector spaces $R_{\leq d}$, which simplifies the implementations to some extent. We use the notation

$$|\alpha|_R := \min\{d \in \mathbb{N} \mid x^\alpha \in R_{\leq d}\}$$

for $\alpha \in \mathbb{N}^n$ to signify the degree with respect to the chosen filtration of R – either total degree or max-degree – so the monomials $\{x^\alpha\}_{|\alpha|_R \leq d}$ form a basis of $R_{\leq d}$.

Remark 4.7.1. Although in principle it is possible to consider $L = R = \mathbb{C}[x_1, \dots, x_n]$ together with an involution defined by $(\sum_\alpha p_\alpha x^\alpha)^\circ = \sum_\alpha \overline{p_\alpha} x^\alpha$ for $\sum_\alpha p_\alpha x^\alpha \in L$, where the involution on \mathbb{C} is complex conjugation, this case is not usually of particular interest. To see this, let $\sigma: L \rightarrow \mathbb{C}$ be a \mathbb{C} -linear map such that the sesquilinear form $\langle -, - \rangle_\sigma$ is positive-semidefinite on $R_{\leq d}$ for some $d \in \mathbb{N}$. Let H be the Gramian with respect to the monomial basis of $R_{\leq d}$. Then H is Hermitian, as it is positive-semidefinite, and it is symmetric since it is a Hankel matrix as discussed in Remark 3.1.7. In particular, this means that H is already a real matrix, so instead we can work with $\mathbb{k} = \mathbb{R}$ in the first place. \diamond

Remark 4.7.2. The results in this section are proved with Corollary 4.5.6 in mind, which forces us to work with positive-semidefinite forms. The positive-semidefiniteness is primarily used as a sufficient criterion for asserting that the constructed matrix pencils are regular, as follows from Lemma 4.2.13. If, in some settings, the regularity of the involved pencils can be established by different means (cf. Section 4.2.2), then it is possible to rephrase the results here in terms of Theorem 4.5.4, which does not assume positive-semidefiniteness. \diamond

Remark 4.7.3. In Algorithm 4.3, we allow the weights $\lambda_{ij} \in \mathbb{C}$, $0 \leq i \leq s, 0 \leq j \leq r$, to be complex numbers, even when $\mathbb{k} = \mathbb{R}$, by replacing the \mathbb{R} -linear maps σ_j by their complexification. If in fact the weights $\lambda_{ij} \in \mathbb{R}$ are all real, then the computations can be performed over \mathbb{R} instead. \diamond

We can now prove the correctness of Algorithm 4.3.

Theorem 4.7.4. *Algorithm 4.3 works and requires the moments*

$$\tau_i((x^\alpha)^\circ x^\beta), \quad |\alpha|_R, |\beta|_R \leq d + \delta,$$

for $0 \leq i \leq s$.

Proof. First, note that it follows from Lemma 1.2.3 that $\bigcap_{k=0, k \neq j}^r (\mathfrak{p}_k \cap R_{\leq d}) \setminus \mathfrak{p}_j$ is not empty for any $0 \leq j \leq r$ if d is sufficiently large. If $L = R$, this is immediately clear since $\bigcup_{d \in \mathbb{N}} R_{\leq d} = R$. Likewise, if L is the Laurent ring, we may restrict to the polynomial ring R and apply the same argument by Lemma 1.2.5.

Algorithm 4.3 Recovery of weights from several weighted sums (degree of generators known)

Input: Natural numbers $s, d, \delta \in \mathbb{N}$ and \mathbb{C} -linear maps τ_i , $0 \leq i \leq s$.

Assumptions: The maps are of the form $\tau_i = \sum_{j=0}^r \lambda_{ij} \sigma_j$ for some $r \leq s$, where

$(\lambda_{ij})_{ij} \in \mathbb{C}^{(s+1) \times (r+1)}$ is a matrix of rank $r+1$ and $\sigma_j: L \rightarrow \mathbb{k}$ are \mathbb{k} -linear maps for which the associated sesquilinear forms $\langle -, - \rangle_{\sigma_j}$ are positive-semidefinite on $R_{\leq d+\delta}$.

Additionally, assume that there exist prime ideals $\mathfrak{p}_j \subseteq L$ with $\mathfrak{p}_j \subseteq \ker \sigma_j$ for $0 \leq j \leq r$ such that the induced sesquilinear form on $R_{\leq d+\delta}/(\mathfrak{p}_j \cap R_{\leq d+\delta})$ associated to σ_j is non-degenerate and such that $\mathfrak{p}_k \not\subseteq \mathfrak{p}_j$ for all $k \neq j$. Moreover, assume that d and δ are sufficiently large such that $\bigcap_{k=0, k \neq j}^r (\mathfrak{p}_k \cap R_{\leq d}) \setminus \mathfrak{p}_j \neq \emptyset$ holds and \mathfrak{p}_j is generated by $\mathfrak{p}_j \cap R_{\leq d}$ for all $0 \leq j \leq r$. Denote $\mathfrak{a} := \mathfrak{p}_0 \cap \dots \cap \mathfrak{p}_r$.

Output: r as well as $[\lambda_{0j} : \dots : \lambda_{sj}]$ for $0 \leq j \leq r$ (up to permutations in j).

1: Set $M_i := (\tau_i((x^\alpha)^\circ x^\beta))_{|\alpha|_R, |\beta|_R \leq d+\delta}$ for every $0 \leq i \leq s$.

2: Compute

$$\bigcap_{i=0}^s \ker M_i = \mathfrak{p}_0 \cap \dots \cap \mathfrak{p}_r \cap R_{\leq d+\delta} = \mathfrak{a} \cap R_{\leq d+\delta}. \quad (4.17)$$

3: For $0 \leq i \leq s$, let M'_i be the Gramian matrix of the induced form associated to τ_i with respect to some basis of $R_{\leq d+\delta}/(\mathfrak{a} \cap R_{\leq d+\delta})$. Then (M'_0, \dots, M'_s) is a regular matrix pencil.

4: Compute the set \mathfrak{V}' of eigenspaces of the pencil (M'_0, \dots, M'_s) .

5: Set $\Gamma := \emptyset$.

6: **for** $V' \in \mathfrak{V}'$ **do**

7: Set $V := V' + \mathfrak{a} \cap R_{\leq d+\delta} \subseteq R_{\leq d+\delta}$.

8: Compute

$$W := \bigcap_{\alpha \in \mathbb{N}^n, |\alpha|_R \leq \delta} \frac{V \cap x^\alpha R_{\leq d}}{x^\alpha} = \{q \in R_{\leq d} \mid qx^\alpha \in V \text{ for } |\alpha|_R \leq \delta\}.$$

9: Set $W' := W/(\mathfrak{a} \cap R_{\leq d+\delta} \cap W)$.

10: If $W' \neq 0$, add the eigenvalue corresponding to V' to the set Γ .

11: Set $r := \#\Gamma - 1$.

12: Return r and $\Gamma = \{[\lambda_{0j} : \dots : \lambda_{sj}] \in \mathbb{P}_{\mathbb{C}}^s \mid 0 \leq j \leq r\}$.

We define the matrices $H_j := \sigma_j((x^\alpha)^\circ x^\beta)|_{\alpha|_R, |\beta|_R \leq d+\delta}$, which are positive-semidefinite by assumption on $\langle -, - \rangle_{\sigma_j}$ and satisfy $M_i = \sum_{j=0}^r \lambda_{ij} H_j$ for $0 \leq i \leq s$. Since the matrix $(\lambda_{ij})_{ij}$ is of full rank, we have

$$\bigcap_{i=0}^s \ker M_i = \bigcap_{j=0}^r \ker H_j.$$

It is clear that $\mathfrak{p}_j \cap R_{\leq d+\delta} \subseteq \ker H_j$ for every $0 \leq j \leq r$ as $\mathfrak{p}_j \subseteq \ker \sigma_j$. Due to the assumption that the induced sesquilinear form on $R_{\leq d+\delta}/(\mathfrak{p}_j \cap R_{\leq d+\delta})$ associated to σ_j is non-degenerate, Lemma 3.1.17 implies in conjunction with Lemma 3.1.16 that, in fact, we have $\mathfrak{p}_j \cap R_{\leq d+\delta} = \ker H_j$ for all j . This establishes the equality (4.17).

Denote by H'_j the matrices corresponding to H_j with respect to a basis of the space $R_{\leq d+\delta}/(\mathfrak{a} \cap R_{\leq d+\delta})$. As these matrices are positive-semidefinite and since $\mathfrak{a} \cap R_{\leq d+\delta}$ is the intersection of the kernels of H_j , the matrices (H'_0, \dots, H'_r) form a regular matrix pencil by Lemma 4.2.13. As the matrix $(\lambda_{ij})_{ij}$ is of rank $r+1$, the pencil (M'_0, \dots, M'_s) is regular as well by Proposition 4.2.20 (1), as claimed in Line 3 of the algorithm.

By hypothesis, we have $\bigcap_{k=0, k \neq j}^r (\mathfrak{p}_k \cap R_{\leq d}) \setminus \mathfrak{p}_j \neq \emptyset$. This means, for every $0 \leq j \leq r$, the regular pencil (H'_0, \dots, H'_r) has a non-trivial eigenspace for the eigenvalue $e_j := [0 : \dots : 0 : 1 : 0 : \dots : 0] \in \mathbb{P}_{\mathbb{k}}^r$ which is non-zero at position j only. Indeed, if $p_j \in \bigcap_{k=0, k \neq j}^r (\mathfrak{p}_k \cap R_{\leq d}) \setminus \mathfrak{p}_j$ and p'_j denotes its residue class in $R_{\leq d+\delta}/(\mathfrak{a} \cap R_{\leq d+\delta})$, then $p_j \not\equiv 0 \pmod{\mathfrak{a} \cap R_{\leq d+\delta}}$ since $p_j \notin \mathfrak{p}_j$, so $p'_j \neq 0$. Furthermore, we have $\sigma_k((x^\alpha)^\circ p_j) = 0$ for all $k \neq j$ and all $|\alpha|_R \leq d+\delta$, since $(x^\alpha)^\circ p_j \in \mathfrak{p}_k \subseteq \ker \sigma_k$. Thus,

$$H'_k p'_j = 0 \cdot H'_j p'_j = 0,$$

so, by Proposition 4.2.5, p'_j is an eigenvector with eigenvalue e_j as claimed. If $\mathbb{k} \neq \mathbb{C}$, we now switch to the complexification of the matrices H'_0, \dots, H'_r as well as the eigenvectors. Then, due to Proposition 4.2.20, the eigenvectors of (H'_0, \dots, H'_r) and (M'_0, \dots, M'_s) are the same. This means that the pencil (M'_0, \dots, M'_s) has at least $r+1$ different eigenvalues and eigenspaces, which are computed in Line 4.

It follows from Corollary 4.5.6 that the non-trivial spaces W' computed in Lines 6 to 10 are subspaces of exactly the $r+1$ eigenspaces referred to above.

Therefore, the subset Γ of the eigenvalues of (M'_0, \dots, M'_s) has cardinality $r+1$ and, by Proposition 4.2.20 (3), consists of the image of e_0, \dots, e_r under the map induced by $(\lambda_{ij})_{ij}$, namely the points

$$[\lambda_{0j} : \dots : \lambda_{sj}] \in \mathbb{P}_{\mathbb{C}}^s, \quad 0 \leq j \leq r, \quad (4.18)$$

as claimed in Line 12. \square

Remark 4.7.5. Algorithm 4.3 and Algorithm 4.4 described below are primarily based on the following idea. The elements contained in the non-empty sets

$$\bigcap_{\substack{k=0 \\ k \neq j}}^r (\mathfrak{p}_k \cap R_{\leq d+\delta}) \setminus \mathfrak{p}_j, \quad 0 \leq j \leq r,$$

are Lagrange-like polynomials in the sense of Remark 4.3.3. They are representatives of eigenvectors of the pencils (M'_0, \dots, M'_s) and (H'_0, \dots, H'_r) (as defined in the proof of Theorem 4.7.4). The eigenvalues of (M'_0, \dots, M'_s) corresponding to these eigenvectors are exactly the values listed in (4.18).

However, in some rare cases, the two pencils can have additional eigenvalues with eigenvectors that are not of this form. For example, this can happen when the varieties corresponding to $\mathfrak{p}_0, \dots, \mathfrak{p}_r$ satisfy particular symmetry conditions, which can be regarded as an exceptional situation (cf. Example 4.3.19). Therefore, we seek to find subspaces of the eigenspaces that have additional structure, in order to guarantee that we are able to detect those eigenvectors that correspond to the eigenvalues (4.18), only.

The main observation is that, for an element $p \in \bigcap_{k=0, k \neq j}^r (\mathfrak{p}_k \cap R_{\leq d}) \setminus \mathfrak{p}_j$, we also have $x^\alpha p \in \bigcap_{k=0, k \neq j}^r \mathfrak{p}_k \cap R_{\leq d+\delta}$ for all $|\alpha|_R \leq \delta$, so p has the property that all its multiples $x^\alpha p$, $|\alpha|_R \leq \delta$, are representatives of eigenvectors in a fixed eigenspace of the pencils, as long as $x^\alpha p \notin \mathfrak{p}_j$. If on the other hand $x^\alpha p \in \mathfrak{p}_j$, then $x^\alpha p \equiv 0 \pmod{\mathfrak{a} \cap R_{\leq d+\delta}}$, so it trivially represents an element of the eigenspace, as well. Now, the results of Section 4.5 show that, for suitable choice of δ , the elements from the sets

$$\bigcap_{\substack{k=0 \\ k \neq j}}^r (\mathfrak{p}_k \cap R_{\leq d}) \setminus \mathfrak{p}_j, \quad 0 \leq j \leq r,$$

are the only eigenvectors with this additional property, which allows us to filter out those eigenspaces of (M'_0, \dots, M'_s) that do not correspond to the eigenvalues listed in (4.18). \diamond

Remark 4.7.6. By Theorem 3.4.11, the assumptions on non-degenerateness in Algorithm 4.3 are satisfied, in particular, when $\sigma_0, \dots, \sigma_r$ are the moment functionals of measures with Zariski-dense support in the varieties corresponding to the ideals $\mathfrak{p}_0, \dots, \mathfrak{p}_r$. \diamond

Remark 4.7.7. Algorithm 4.3 recovers the weights $[\lambda_{0j} : \dots : \lambda_{sj}]$ up to scaling and permutation in j . From this, we can recover the moment matrices

$$H_j := \sigma_j((x^\alpha)^\circ x^\beta)_{|\alpha|_R, |\beta|_R \leq D},$$

for some $D \in \mathbb{N}$, corresponding to the individual components, up to scaling, by solving the linear system

$$M_i = \sum_{j=0}^r \lambda_{ij} H_j, \quad 0 \leq i \leq s, \quad (4.19)$$

as the matrix $(\lambda_{ij})_{ij}$ has full column rank. Thus, when $\mathfrak{p}_j \cap R_{\leq D} = \ker H_j$ we obtain the generators of \mathfrak{p}_j of degree at most D . As explained in the proof of Theorem 4.7.4, the requirement $\mathfrak{p}_j \cap R_{\leq D} = \ker H_j$ is satisfied when the induced sesquilinear form on $R_{\leq D}/(\mathfrak{p}_j \cap R_{\leq D})$ is non-degenerate. Due to the assumption of positive-semidefiniteness, this holds for all D that are chosen to be smaller than the concrete bound $d + \delta$ in the algorithm.

An alternative approach for computing $\mathfrak{p}_j \cap R_{\leq \delta}$ for $\delta \in \mathbb{N}$ is based on the following observation. Assume that $p \in \bigcap_{k=0, k \neq j}^r \mathfrak{p}_k \cap R_{\leq d}$, so p is an eigenvector as computed in Line 8. Without loss of generality, we can assume that $p \notin \mathfrak{p}_j$. Additionally, let $q \in R_{\leq \delta}$ such that $\sigma_j((x^\alpha)^\circ qp) = 0$ for all $|\alpha|_R \leq d + \delta$. As σ_j induces a non-degenerate sesquilinear form on $R_{\leq d+\delta}/(\mathfrak{p}_j \cap R_{\leq d+\delta})$, this implies that $qp \in \mathfrak{p}_j$ and therefore $q \in \mathfrak{p}_j$. Conversely, if $q \in \mathfrak{p}_j \cap R_{\leq \delta}$, then also $\sigma_j((x^\alpha)^\circ qp) = 0$ holds, due to $(x^\alpha)^\circ qp \in \mathfrak{p}_j \subseteq \ker \sigma_j$, so the two conditions on q are equivalent.

Note further that $M_i(qp) = (\lambda_{ij} \sigma_j((x^\alpha)^\circ qp))|_{|\alpha|_R \leq d+\delta}$, since $(x^\alpha)^\circ qp \in \mathfrak{p}_k \subseteq \ker \sigma_k$ for all $0 \leq k \leq r$ with $k \neq j$. As the matrix $(\lambda_{ij})_{ij}$ has full column rank, there exists an index i for which $\lambda_{ij} \neq 0$, so it follows from the above characterization that $q \in \mathfrak{p}_j \cap R_{\leq \delta}$ if and only if $M_i(qp) = 0$ for all i . In summary, starting from the eigenvectors computed in Line 8, we may compute the elements $q \in R_{\leq \delta}$ with this property to obtain generators of $\mathfrak{p}_j \cap R_{\leq \delta}$. \diamond

In particular special cases, there can be other algorithms that directly recover the components \mathfrak{p}_j , $0 \leq j \leq r$, instead of solving the linear system (4.19). For instance, the multi-snapshot ESPRIT algorithm deals with the case in which the ideals \mathfrak{p}_j correspond to single points in dimension $n = 1$; see [LZGL21]. Despite not solving the linear system (4.19), their algorithm requires that the weighting matrix $(\lambda_{ij})_{ij}$ has full column rank as well, so this seems to be a natural assumption.

Remark 4.7.8. The argument in Remark 4.7.7 only involves a single polynomial $p \in \bigcap_{k=0, k \neq j}^r \mathfrak{p}_k \cap R_{\leq d}$. As the property $\sigma_j((x^\alpha)^\circ qp) = 0$ holds for $q \in \mathfrak{p}_j \cap R_{\leq \delta}$ and for arbitrary $p \in \bigcap_{k=0, k \neq j}^r \mathfrak{p}_k \cap R_{\leq d}$, we would like to decrease the necessary degree bound on α by quantifying over all p (possibly to $|\alpha|_R \leq \delta$ if \mathfrak{p}_j is generated by $\mathfrak{p}_j \cap R_{\leq \delta}$). This problem seems difficult to approach, so we leave it for further investigation. \diamond

If an upper bound on the degrees of generators of the involved ideals is not known, then Algorithm 4.3 is not applicable. Instead, recovery is possible by Algorithm 4.4, as is proved in the following. The main difference from the previous algorithm is that the new one includes a loop which successively increments the degree δ until it is large enough for reconstruction.

Theorem 4.7.9. *Algorithm 4.4 works and requires finitely many moments of τ_0, \dots, τ_s . If $D \in \mathbb{N}$ is the smallest number such that the ideals \mathfrak{p}_j are generated by $\mathfrak{p}_j \cap R_{\leq D}$ for all $0 \leq j \leq r$, then the moments at $(x^\alpha)^\circ x^\beta$ for $|\alpha|_R, |\beta|_R \leq d + D$ are sufficient.*

Proof. The proof is similar to that of Theorem 4.7.4. Observe that, if d is sufficiently large, then

$$\bigcap_{\substack{k=0 \\ k \neq j}}^r (\mathfrak{p}_k \cap R_{\leq d}) \setminus \mathfrak{p}_j$$

Algorithm 4.4 Recovery of weights from several weighted sums (degree of generators unknown)

Input: Natural numbers $s, d \in \mathbb{N}$ and \mathbb{C} -linear maps τ_i , $0 \leq i \leq s$.

Assumptions: The maps are of the form $\tau_i = \sum_{j=0}^r \lambda_{ij} \sigma_j$, for some $r \leq s$, where $(\lambda_{ij})_{ij} \in \mathbb{C}^{(s+1) \times (r+1)}$ is a matrix of rank $r+1$ and $\sigma_j: L \rightarrow \mathbb{k}$ are \mathbb{k} -linear maps for which the associated sesquilinear forms $\langle -, - \rangle_{\sigma_j}$ are positive-semidefinite on R . Additionally, assume that there exist prime ideals $\mathfrak{p}_j \subseteq L$ with $\mathfrak{p}_j \subseteq \ker \sigma_j$ for $0 \leq j \leq r$ such that the induced sesquilinear form on $R_{\leq d+\delta}/(\mathfrak{p}_j \cap R_{\leq d+\delta})$ associated to σ_j is non-degenerate for $\delta = 0, 1, 2, \dots$ and such that $\mathfrak{p}_k \not\subseteq \mathfrak{p}_j$ for all $k \neq j$. Moreover, assume that d is sufficiently large such that $\bigcap_{k=0, k \neq j}^r (\mathfrak{p}_k \cap R_{\leq d}) \setminus \mathfrak{p}_j \neq \emptyset$ for all $0 \leq j \leq r$. Denote $\mathfrak{a} := \mathfrak{p}_0 \cap \dots \cap \mathfrak{p}_r$ as well as $M_{i,e} := (\tau_i((x^\alpha)^\circ x^\beta))_{|\alpha|_R, |\beta|_R \leq e}$ for every $0 \leq i \leq r$ and $e \in \mathbb{N}$.

Output: r as well as $[\lambda_{0j} : \dots : \lambda_{sj}]$ for $0 \leq j \leq r$ (up to permutations in j).

- 1: Set $\delta := 0$.
- 2: Compute $\bigcap_{i=0}^s \ker M_{i,d+\delta} = \mathfrak{p}_0 \cap \dots \cap \mathfrak{p}_r \cap R_{\leq d+\delta} = \mathfrak{a} \cap R_{\leq d+\delta}$.
- 3: For $0 \leq i \leq s$, let M'_i be the Gramian matrix of the induced form associated to τ_i with respect to some basis of $R_{\leq d+\delta}/(\mathfrak{a} \cap R_{\leq d+\delta})$. Then (M'_0, \dots, M'_s) is a regular matrix pencil.
- 4: Compute the set \mathfrak{V}' of eigenspaces of the pencil (M'_0, \dots, M'_s) . Denote the set of corresponding eigenvalues by $\Gamma \subseteq \mathbb{P}_{\mathbb{C}}^s$.
- 5: **if** $\delta = 0$ **then**
- 6: Let r be the dimension of the projective subspace spanned by Γ .
- 7: **if** $\#\Gamma > r + 1$ and $\delta \geq 1$ **then**
- 8: **for** $V' \in \mathfrak{V}'$ **do**
- 9: Set $V := V' + \mathfrak{a} \cap R_{\leq d+\delta} \subseteq R_{\leq d+\delta}$.
- 10: Compute

$$W := \bigcap_{\alpha \in \mathbb{N}^n, |\alpha|_R \leq \delta} \frac{V \cap x^\alpha R_{\leq d}}{x^\alpha} = \{q \in R_{\leq d} \mid qx^\alpha \in V \text{ for } |\alpha|_R \leq \delta\}.$$

- 11: Set $W' := W/(\mathfrak{a} \cap R_{\leq d+\delta} \cap W)$.
 - 12: If $W' = 0$, remove the eigenvalue corresponding to V' from the set Γ .
 - 13: If Γ consists of more than $r + 1$ eigenvalues, increment δ by 1 and go to Line 2.
 - 14: Return $\Gamma = \{[\lambda_{0j} : \dots : \lambda_{sj}] \in \mathbb{P}_{\mathbb{C}}^s \mid 0 \leq j \leq r\}$.
-

is non-empty for all $0 \leq j \leq r$ by Lemma 1.2.3. In that case, also

$$\bigcap_{\substack{k=0 \\ k \neq j}}^r (\mathfrak{p}_k \cap R_{\leq d+\delta}) \setminus \mathfrak{p}_j$$

is non-empty for all $0 \leq j \leq r$ and all $\delta \in \mathbb{N}$.

Following the arguments in the proof of Theorem 4.7.4, we infer that \mathfrak{W}' , computed in Line 4, consists of at least $r + 1$ eigenspaces, for any $\delta \in \mathbb{N}$. Every eigenvalue of these eigenspaces is contained in the image of the injective map $\mathbb{P}_{\mathbb{C}}^r \rightarrow \mathbb{P}_{\mathbb{C}}^s$ induced by $(\lambda_{ij})_{ij}$. As at least the $r + 1$ distinct elements $[\lambda_{0j} : \cdots : \lambda_{sj}]$, $0 \leq j \leq r$, occur as eigenvalues, this implies that r must be equal to the dimension of the projective subspace spanned by Γ , in Line 6. Thus, in Line 7, the cardinality of Γ is at least $r + 1$ and the computation can finish if the cardinality is exactly $r + 1$.

Otherwise, by the computation in Lines 7 to 12, we filter out undesirable eigenvalues that do not satisfy the property described in Remark 4.7.5. The eigenvalues $[\lambda_{0j} : \cdots : \lambda_{sj}]$, $0 \leq j \leq r$, are retained by this operation. When $\delta = 0$, this computation is redundant and can be skipped, so we only carry it out if $\delta \geq 1$.

If in Line 13 the set of eigenvalues still has cardinality larger than $r + 1$, we start over and repeat the computations for $\delta + 1$. As each space W computed in Line 10 in iteration $\delta + 1$ is contained in one of the spaces computed in the previous iteration and since $\mathfrak{a} \cap R_{\leq d+\delta} \subseteq \mathfrak{a} \cap R_{\leq d+\delta+1}$, it is clear that the cardinality of Γ in Line 13 does not grow when δ increases. However, if δ is sufficiently large such that all the ideals \mathfrak{p}_j , $0 \leq j \leq r$, are generated by elements in $\mathfrak{p}_j \cap R_{\leq \delta}$, then Γ has cardinality exactly $r + 1$ by Theorem 4.7.4. Thus, the cardinality of Γ in Line 13 is a weakly monotonically decreasing sequence in δ that stagnates at the value $r + 1$, which shows that the algorithm terminates. \square

Remark 4.7.10. In practice, we expect that Algorithm 4.4 already terminates for small δ , usually $\delta = 0$ or in some cases $\delta = 1$, so the computation does not commonly need all the moments up to the upper bound given in Theorem 4.7.9. \diamond

Remark 4.7.11. In Algorithm 4.4, we require that, for $0 \leq j \leq r$, the induced sesquilinear form on $R_{\leq d+\delta}/(\mathfrak{p}_j \cap R_{\leq d+\delta})$ associated to σ_j is non-degenerate for all $\delta \in \mathbb{N}$. This condition only needs to be satisfied for the values of δ that actually arise in the computation, so, in the worst case, it is sufficient if this criterion is met for all $\delta \in \{0, \dots, D\}$, where D is chosen as in Theorem 4.7.9.

In practice, such a condition may be difficult to ascertain, but sometimes this will be clear by other information about an application, for example in case of non-negative measures that are densely supported on irreducible varieties with respect to the Zariski topology. \diamond

Remark 4.7.12. For simplicity of exposition, Algorithms 4.3 and 4.4 assume that the underlying ideals are prime. More generally, we can extend this to the case of ideals that are not necessarily prime, but satisfy the assumptions of Theorem 4.5.4, most notably

property (4) of Lemma 4.5.2. From a geometric point of view, this corresponds to components that are unions of irreducible varieties, so each component is a possibly reducible variety. More precisely, the $r + 1$ components are defined by ideals $\mathfrak{a}_0, \dots, \mathfrak{a}_r$ of the form

$$\mathfrak{a}_j = \mathfrak{p}_{j,1} \cap \dots \cap \mathfrak{p}_{j,m_j},$$

where $\mathfrak{p}_{j,k}$ are prime ideals for $0 \leq j \leq r$ and $1 \leq k \leq m_j$. Additionally, the components must not overlap, meaning that $\mathfrak{p}_{j,k} \not\supseteq \mathfrak{p}_{j',k'}$ whenever $j \neq j'$ or $k \neq k'$. The components may be of different dimensions, though.

In a sense, the bound on δ prescribed by condition (4) of Lemma 4.5.2 is larger than in the case of irreducible components, condition (3): for irreducible components, δ is bounded by the maximum degree of generating polynomials of each component; in the case of reducible components, the bound can be almost twice as large. However, in Algorithm 4.4, the theoretical bound on δ is merely a sufficient criterion that is not attained in practice, as usually small values like $\delta = 0$ or $\delta = 1$ seem to be enough already. Therefore, we expect that it does not play a major role that, in the reducible case, the bound on δ can be almost twice as large as in the irreducible case. \diamond

Remark 4.7.13. Another potential improvement of Algorithm 4.4 involves the computation of eigenspaces and eigenvalues. In each iteration of the algorithm, all the eigenvalues and eigenspaces of the pencil are recomputed. Depending on the concrete implementation for this computation, it may be possible to avoid some of this computational effort by taking into account the eigenvalues from the previous iteration. Only those eigenvalues that are contained in the set Γ in Line 13 are relevant for the computation, so, for efficiency, the following iteration may omit the computation of eigenspaces that do not correspond to eigenvalues in Γ .

Note that, in general, it seems difficult to determine a suitable degree d , a priori. If d is chosen too small, one might find fewer than $r + 1$ eigenspaces, so that reconstruction is not possible using the algorithm. \diamond

4.8 Numerical implementation

This section addresses details that are relevant for an implementation of the algorithms developed in the preceding sections. This concerns the computation of intersections of vector spaces as well as the computation of eigenvalues of a pencil. We finish with a demonstration of the algorithms on some examples, in Section 4.8.3.

4.8.1 Intersection of vector subspaces

A crucial ingredient for Algorithms 4.3 and 4.4 is the computation of intersections of vector subspaces. Here we briefly discuss how this can be implemented.

Definition 4.8.1. Assume that V denotes a complex matrix with linearly independent columns. Then

$$\pi_V := V(V^*V)^{-1}V^*$$

is the orthogonal *projection matrix* onto the column space of V ; see [Mey00, Chapter 5.13]. Indeed, its image is contained in the column space of V , it satisfies $\pi_V V = V$ and $\pi_V w = 0$ for any vector w in the orthogonal complement of the column space of V .

Lemma 4.8.2. *If V_1, \dots, V_r are complex matrices with linearly independent columns each, then a vector v is contained in the intersection of their column spaces if and only if*

$$\pi_{V_1} \cdots \pi_{V_r} v = v.$$

Otherwise, it holds that

$$\|\pi_{V_1} \cdots \pi_{V_r} v\| < \|v\|.$$

Proof. As π_{V_j} is a projection matrix for $1 \leq j \leq r$, we have $\|\pi_{V_j} v\| \leq \|v\|$ for any vector v and $\|\pi_{V_j} v\| < \|v\|$ if and only if v is not contained in the column space of V_j , as a consequence of the triangular inequality. Therefore, it follows that $\|\pi_{V_1} \cdots \pi_{V_r} v\| < \|v\|$ if v is not contained in some of the column spaces. If, on the other hand, v is contained in the intersection of the column spaces of V_1, \dots, V_r , then we have $\pi_{V_1} \cdots \pi_{V_r} v = v$. \square

This allows us to compute the intersection of vector subspaces, numerically.

Lemma 4.8.3. *Let V_1, \dots, V_r be complex matrices with linearly independent columns. Then the intersection of their column spaces is spanned by those right singular vectors of the matrix $\pi_{V_1} \cdots \pi_{V_r}$ for which the corresponding singular values are equal to 1.*

Proof. Assume that $\pi_{V_1} \cdots \pi_{V_r} = \sum_{i=1}^n u_i \varsigma_i v_i^*$ is the singular value decomposition with singular values $\varsigma_i \geq 0$ and left and right singular vectors u_i, v_i , $1 \leq i \leq n$, respectively.

First, observe that all the singular values are contained in the interval $[0, 1]$. Indeed, assuming that the largest singular value ς_1 is greater than 1, then, for the corresponding left and right singular vectors u_1 and v_1 , it follows that

$$\|\pi_{V_1} \cdots \pi_{V_r} v_1\| = \|u_1 \varsigma_1 v_1^* v_1\| = \varsigma_1 > 1 = \|v_1\|,$$

which is a contradiction to Lemma 4.8.2.

In order to prove the statement, we consider an arbitrary linear combination of the right singular vectors $v = \sum_{i=1}^n c_i v_i$ with $c_i \in \mathbb{C}$, $1 \leq i \leq n$, for which we obtain

$$\|\pi_{V_1} \cdots \pi_{V_r} v\|^2 = \left\| \sum_{i=1}^n u_i \varsigma_i v_i^* \sum_{k=1}^n c_k v_k \right\|^2 = \left\| \sum_{i=1}^n u_i \varsigma_i c_i \right\|^2 = \sum_{i=1}^n \varsigma_i^2 |c_i|^2.$$

On the other hand, we have

$$\|v\|^2 = \sum_{i=1}^n |c_i|^2.$$

By Lemma 4.8.2, the vector v is contained in the intersection of the column spaces of V_1, \dots, V_r if and only if these two expressions are equal. As all the singular values are less than or equal to 1, this is the case if and only if, for every $1 \leq i \leq n$, either $\varsigma_i = 1$ or $c_i = 0$, which proves the claim. \square

Remark 4.8.4 (Intersection of column spaces). Thus, by computing the singular vectors of the matrix $\pi_{V_1} \cdots \pi_{V_r}$ with singular values equal to (or numerically close to) 1, we obtain an orthonormal basis of the intersection of the column spaces of V_1, \dots, V_r ; see also [GV96, Chapter 12.4.4].

In practice, we first choose orthonormal bases for the column spaces of V_1, \dots, V_r , for example by QR-decomposition. If Q_1, \dots, Q_r are matrices with orthonormal columns and that have the same column spaces as those of V_1, \dots, V_r , respectively, then we merely need to compute a singular value decomposition of

$$\pi_{V_1} \cdots \pi_{V_r} = Q_1 Q_1^* \cdots Q_r Q_r^*, \quad (4.20)$$

since we have $\pi_{V_j} = V_j V_j^\dagger = Q_j Q_j^*$ for $1 \leq j \leq r$ (cf. [Mey00, Equation (5.13.4)]).

A small optimization of this procedure consists of computing the singular value decomposition of the matrix

$$Q_1^* Q_2 Q_2^* \cdots Q_{r-1} Q_{r-1}^* Q_r \quad (4.21)$$

instead of (4.20) (cf. [GV96, Algorithm 12.4.3]). This is made precise in Lemma 4.8.5. Compared to the approach outlined above, this avoids one matrix multiplication by Q_1 and only requires the singular value decomposition of the matrix (4.21), which has smaller size than (4.20). To this end, it seems economical to choose an order for the matrices Q_1, \dots, Q_r under which Q_1 and Q_r are matrices with few columns, so that only the singular value decomposition of a small matrix is computed. \diamond

Lemma 4.8.5. *Let Q_1, \dots, Q_r be complex matrices with orthonormal columns and let v_1, \dots, v_s be those right singular vectors of the matrix (4.21) for which the corresponding singular values are equal to 1. Then the vectors $Q_r v_1, \dots, Q_r v_s$ form an orthonormal basis of the intersection of the column spaces of Q_1, \dots, Q_r .*

Proof. Assume that $\sum_{i=1}^n u_i \varsigma_i v_i^*$ is the singular value decomposition of (4.21) with singular values $\varsigma_i \geq 0$ and left and right singular vectors u_i, v_i , $1 \leq i \leq n$, respectively, where $s \leq n$. Then we claim that

$$Q_1 Q_1^* \cdots Q_r Q_r^* = \sum_{i=1}^n (Q_1 u_i) \varsigma_i (Q_r v_i)^*$$

is a singular value decomposition. Indeed, we have

$$(Q_r v_i)^* Q_r v_j = v_i^* v_j = \delta_{ij}$$

for $1 \leq i, j \leq n$, so the vectors $Q_r v_1, \dots, Q_r v_n$ and, similarly, $Q_1 u_1, \dots, Q_1 u_n$ are orthonormal. Thus, the statement follows from Lemma 4.8.3 and the equality (4.20). \square

If we are interested in computing the intersection of the kernels of complex matrices A_1, \dots, A_r (rather than the intersection of their column spaces), one way to accomplish this is by first computing an orthonormal basis of $\ker A_j$, for each $1 \leq j \leq r$, and then computing the intersection $\bigcap_{j=1}^r \ker A_j$, as outlined above in Remark 4.8.4 and Lemma 4.8.5. A more direct approach for computing the intersection of matrix kernels employs the following basic fact from linear algebra.

Lemma 4.8.6. *Let $Z \in \mathbb{C}^{n \times m}$ be a matrix with orthonormal columns and let $A \in \mathbb{C}^{s \times n}$. If the columns of a matrix $W \in \mathbb{C}^{m \times q}$ form an orthonormal basis of $\ker(AZ)$, then the columns of ZW are an orthonormal basis of $\text{im}(Z) \cap \ker(A)$.*

The proof is essentially the same as that of [GV96, Theorem 12.4.1], which we provide here for completeness.

Proof. Clearly, $(ZW)^*ZW$ is the identity matrix, so the matrix ZW has orthonormal columns. Moreover, we have $\text{im}(ZW) \subseteq \text{im}(Z)$ and $AZW = 0$, so $\text{im}(ZW) \subseteq \text{im}(Z) \cap \ker(A)$. For the converse, let $x \in \text{im}(Z) \cap \ker(A)$. Then there exists a $y \in \mathbb{C}^m$ with $Zy = x$. Due to $0 = Ax = AZy$, this means that $y \in \ker(AZ)$ and thus there exists a $z \in \mathbb{C}^q$ with $Wz = y$. Therefore, we have $x = ZWz \in \text{im}(ZW)$. \square

Remark 4.8.7 (Intersection of kernels). Lemma 4.8.6 allows us to compute the intersection of the kernels of complex matrices A_1, \dots, A_r that have the same number of columns. First, we compute a singular value decomposition to obtain an orthonormal basis of $\ker A_1$. Denote by Z_1 the matrix whose column vectors are this orthonormal basis, so $\text{im} Z_1 = \ker A_1$. Then, repeatedly apply Lemma 4.8.6 to the matrices Z_{j-1} and A_j for $2 \leq j \leq r$ in order to construct a matrix Z_j whose columns form an orthonormal basis of $\bigcap_{k=1}^j \ker A_k$. Thus, in particular, the columns of Z_r are an orthonormal basis of $\bigcap_{k=1}^r \ker A_k$. This approach is a straightforward generalization of [GV96, Algorithm 12.4.2] to possibly more than two matrices. \diamond

4.8.2 Computation of eigenvalues of a pencil $(\Delta_0, \dots, \Delta_r)$

A fundamental part of the implementation is the computation of generalized eigenvalues of a regular matrix pencil. For this, we use two approaches, which we outline below.

Remark 4.8.8. A quick approach for computing the eigenvalues of a regular matrix pencil $(\Delta_0, \dots, \Delta_r)$ under some genericity assumptions is based on the following. Assume that Δ is a linear combination of $\Delta_0, \dots, \Delta_r$ which is invertible. For instance, we can choose a random linear combination if the pencil is regular, since by (4.2) the set of linear combinations that are singular is a Zariski-closed proper subset and, hence, has measure zero. Then we compute a QZ-decomposition, the generalized Schur form, of the pencil (Δ_0, Δ) such that

$$Q^*(\Delta_0, \Delta)Z = (T_0, T), \quad (4.22)$$

where Q, Z are unitary and T_0, T are upper triangular matrices. The ratios between the diagonal entries of T_0, T are the eigenvalues of (Δ_0, Δ) . We denote these eigenvalues

by $\gamma_1, \dots, \gamma_m \in \mathbb{P}^1$. (Note that these contain the 0-th coordinates of the eigenvalues of the full pencil $(\Delta_0, \dots, \Delta_r)$.) For each of these eigenvalues, we estimate the multiplicity and then compute the corresponding right eigenspaces. For a fixed eigenvalue, say γ_1 with multiplicity $t_1 \in \mathbb{N}$, this can be accomplished by rearranging (4.22) (for example by Givens rotations) such that the t_1 leading diagonal entries of T_0, T correspond exactly to the multiple eigenvalue γ_1 . Then the first t_1 columns of Z span the right eigenspace of γ_1 . Denote this submatrix by Z_1 .

Generically, there is at most one eigenvalue of $(\Delta_0, \dots, \Delta_r)$ that has γ_1 as 0-th coordinate with respect to Δ . If there exists such an eigenvalue of $(\Delta_0, \dots, \Delta_r)$, the corresponding eigenvector must be contained in the span of Z_1 , so we can test whether or not such an eigenvalue exists for the coordinate γ_1 . If so, we can compute the remaining coordinates of the eigenvalue by considering the pencils (Δ_j, Δ) for $1 \leq j \leq r$.

There are some drawbacks to this approach. If the matrices are not generic, then several different eigenvalues of $(\Delta_0, \dots, \Delta_r)$ can have the same 0-th coordinate, which then cannot be distinguished directly. However, it is possible to refine and iterate this by considering $\Delta_1, \Delta_2, \dots$ in order to separate those eigenvalues that are equal in the 0-th coordinate. For more details on this, see for example [HKP04, Section 2]. Alternatively, one may replace the pencil by one that is perturbed by a random linear transformation, which separates all the coordinates of distinct eigenvalues, and obtain the eigenvalues of the original pencil based on Proposition 4.2.20. This mimics a commonly used technique for computing the joint eigenvectors of simultaneously diagonalizable matrices; see for instance [LR12, Section 2] or [Mou18, Corollary 3.1].

Another drawback is a numerical issue. It may happen that two eigenvalues of the pencil $(\Delta_0, \dots, \Delta_r)$ have 0-th coordinates that are very close to each other, so that these eigenvalues are difficult to distinguish by considering just the 0-th coordinate. In this case, it may be more appropriate to first consider another coordinate – however, one cannot detect this a priori and, additionally, other eigenvalues might be close in terms of the other coordinates. This motivates a refined procedure described in Remark 4.8.9 below. \diamond

Remark 4.8.9. In order to overcome the numerical problem that results from projecting onto a single coordinate when computing the eigenvalues of a regular pencil $(\Delta_0, \dots, \Delta_r)$, we employ the following variant of the approach described in Remark 4.8.8. A similar strategy is described in [CPS12], for the case of jointly diagonalizable and, thus, commuting matrices.

Instead of computing each eigenvalue of (Δ_0, Δ) with multiplicity (which is difficult in the presence of numerical noise if two eigenvalues are close), we split the set of eigenvalues into two well-separated clusters. If this is not possible, we consider a different coordinate instead. For each cluster, we compute the *deflating subspaces* containing the corresponding eigenspaces (see e. g. [Dem00, Chapter 2.6.2]), which is numerically more stable than computing all the individual eigenspaces, since the two clusters are well-separated. In this case, the deflating subspaces are insensitive to small perturbations of the matrix

entries, whereas the individual eigenspaces can be very sensitive to small perturbations.

This is accomplished by computing an ordered QZ-decomposition as in (4.22) in such a way that the leading diagonal entries of the triangular matrices correspond to the eigenvalues of the cluster, in which case the leading columns of Q, Z span the left and right deflating subspaces. Denote the submatrices spanning the deflating subspaces by Q_1, Z_1 . Then this allows to reduce to a smaller problem by considering the pencil

$$Q_1^*(\Delta_0, \dots, \Delta_r)Z_1. \quad (4.23)$$

Indeed, if v is an eigenvector of $(\Delta_0, \dots, \Delta_r)$ with eigenvalue γ such that $\Delta_j v = \gamma_j \Delta v$ for $0 \leq j \leq r$ and such that γ_0 belongs to the cluster, then v is contained in the span of Z_1 . Thus, we can write $v = Z_1 w$ for some vector w and it follows that $Q_1^* \Delta_j Z_1 w = \gamma_j Q_1^* \Delta Z_1 w$ for $0 \leq j \leq r$. Note that $Q_1^* \Delta Z_1$ is an upper triangular matrix with non-zero entries on the diagonal, so in particular it is regular. Hence, w is an eigenvector with the same eigenvalue γ of the smaller matrix pencil (4.23).

Rather than further subdividing the cluster with respect to the 0-th coordinate, we continue with the next coordinate, as this makes it more likely to find well-separated clusters of eigenvalues. (It would also be possible to optimize this a bit by choosing the coordinate in which the best-separated two clusters exist, but this comes at the cost of computing multiple additional eigenvalues, so we do not pursue this path.) Repeating this process until no two well-separated clusters of eigenvalues can be found in any coordinate eventually yields (at most) a single eigenvalue of $(\Delta_0, \dots, \Delta_r)$ together with its multiplicity, for each sequence of clusters considered.

To determine whether two eigenvalues $\gamma, \gamma' \in \mathbb{P}_{\mathbb{C}}^1$ of a pencil of two matrices are close to each other, we use the *chordal metric* (cf. [Dem00, Chapter 2.6.5]), which is induced by the identification of $\mathbb{P}_{\mathbb{C}}^1$ with the two-dimensional sphere. It is defined by the quantity

$$|\gamma_0 \gamma'_1 - \gamma_1 \gamma'_0|,$$

if the coordinates of the eigenvalues are normalized such that $|\gamma_0|^2 + |\gamma_1|^2 = 1 = |\gamma'_0|^2 + |\gamma'_1|^2$. \diamond

Remark 4.8.10. To be able to detect the generalized eigenvalues, it is important to choose a suitable tolerance threshold. This presents us with a numerical dilemma: If the tolerance is too large, two eigenvalues that are close but different may be detected as a single eigenvalue; if it is too small, an eigenvalue with multiplicity may be detected as several distinct ones or it may not be detected as an eigenvalue at all. Therefore, if the errors in the input data are too large, we cannot expect to correctly identify the desired eigenvalues, but, if the errors are sufficiently small, then computing the eigenvalues is possible in practice. \diamond

4.8.3 Numerical examples

Here we demonstrate the recovery algorithms developed in this chapter in terms of concrete numerical examples. The computer code for these and several other examples is available in the repository [Wag21], including examples for Algorithms 4.2 to 4.4. The implementation uses the software SageMath [Sag21]; many of the numeric computations are performed by NumPy [Har+20], SciPy [Vir+20] as well as Arb [Joh17]; the eigenvalue computations are performed by LAPACK [And+99].

Example 4.8.11. We consider an example on the two-dimensional torus that involves several algebraic components that are generated by polynomials of max-degree 2. For the measures, we choose the uniform measures μ_0, \dots, μ_r supported on these trigonometric algebraic varieties on the torus, where $r + 1$ is the number of components. We compute to machine precision the moments of several signed measures $\nu_i := \sum_{j=0}^r \lambda_{ij} \mu_j$, for $0 \leq i \leq r$, which we use as input to Algorithm 4.4. Here, the weights $\lambda_{ij} \in \mathbb{C}$ are random complex numbers. To compute the moments, we locally parametrize the curves, which allows us to compute the desired integrals numerically.

As there are no particular symmetries between the components, we expect that the algorithm succeeds already in the first iteration with $\delta = 0$. Moreover, since the components are generated by polynomials of degree 2, it is sufficient to choose the parameter d as $d = 2r$ in order to satisfy the requirement $\bigcap_{k=0, k \neq j}^r (\mathfrak{p}_k \cap R_{\leq d}) \setminus \mathfrak{p}_j \neq \emptyset$ for $0 \leq j \leq r$.

For the tolerance threshold parameter described in Remark 4.8.10, we choose a value of 10^{-8} . Then, up to small numerical errors, the algorithm successfully finds the correct eigenvalues and recovers the weights $[\lambda_{0j} : \dots : \lambda_{sj}]$, $0 \leq j \leq r$, up to scaling. From there, we reconstruct the moment matrices associated to μ_0, \dots, μ_r , again up to scaling, by solving a linear system as described in Remark 4.7.7, which allows us to infer information about the individual components.

Figure 4.2 illustrates the successful reconstruction in case of 3 components; Figures 4.3 and 4.4 show the case of 6 components. For the purpose of visualization, we display $Q_{d,\varepsilon}$, defined in Definition 3.5.4, even in the case of data that is not positive-semidefinite (cf. Remark 3.5.8). For small $\varepsilon > 0$, this produces images that are better localized at the varieties than what is obtained by using the functions P_d or $P_{d,1}$ as approximations, as noted in Example 3.5.15.

Observe that, with $d = 2r$, the degree is not large enough for the existence of non-zero polynomials that vanish on all of the components. In this case, the function $P_{d,1}$ is constant and does not provide any information, at all. This is a major contributing factor to the blurriness of the subfigures visualizing the input data in terms of $Q_{d,\varepsilon}$, which is especially pronounced in Figure 4.2. Although some features of the varieties are recognizable beforehand, the details only become visible after reconstruction, as displayed in the subfigures in the center and bottom row of Figure 4.2 and in Figure 4.4. Finally, note that some of the irreducible algebraic components shown in the figures consist of two disconnected parts on the torus. \diamond




Figure 4.2: The top row displays $Q_{d,\varepsilon}$ for $d = 4$, $\varepsilon = 0.01$ (as defined in Definition 3.5.4), associated to the measures ν_0, ν_1, ν_2 from Example 4.8.11, supported on three trigonometric algebraic curves on \mathbb{T}^2 ; the center row shows $Q_{d,\varepsilon}$ associated to each of the three reconstructed measures, normalized to equal heights; the bottom row combines the reconstructions into a single image.




Figure 4.3: The function $Q_{d,\epsilon}$ for $d = 10$, $\epsilon = 0.01$, associated to the measures ν_0, \dots, ν_5 from Example 4.8.11, which are supported on six trigonometric algebraic curves on \mathbb{T}^2 and are used as input for Algorithm 4.4; the reconstruction is shown in Figure 4.4.

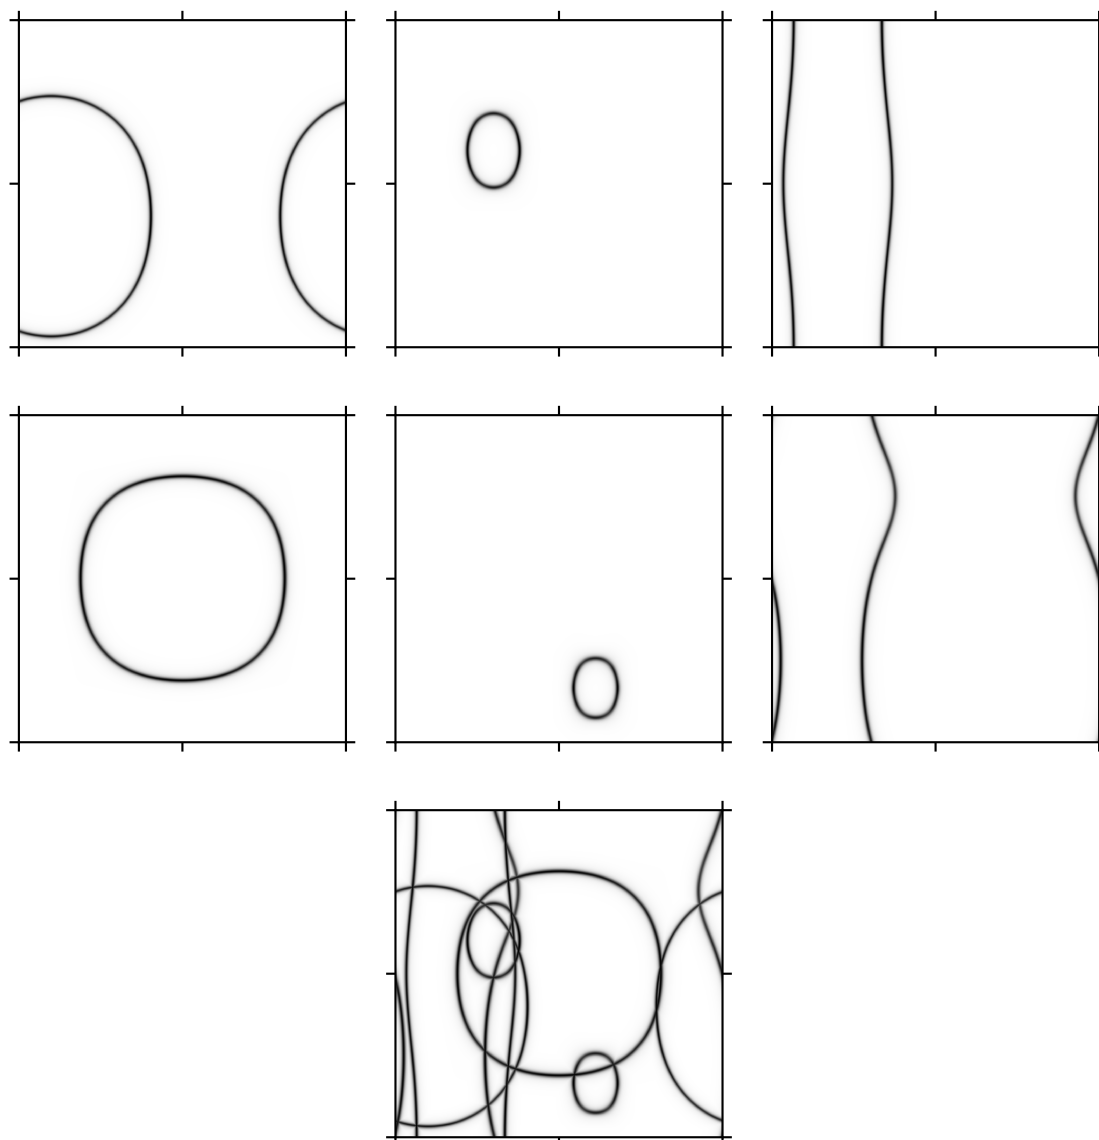


Figure 4.4: The function $Q_{d,\epsilon}$ for $d = 10$, $\epsilon = 0.01$, associated to each of the six measures on \mathbb{T}^2 that are reconstructed from the ones displayed in Figure 4.3, normalized to equal heights; the bottom picture combines all the recovered components into a single image.

As eigenvalues can be computed symbolically for small matrices, it is possible to implement a symbolic variant of the algorithms. We use that to obtain exact results in the following small example.

Example 4.8.12. Let us revisit Example 4.3.19 more explicitly. We consider three circles in the affine plane whose center points lie on a single line through the origin. As explained in detail in Example 4.3.19, it is then expected that Algorithm 4.1 or Algorithm 4.3 for $\delta = 0$ are not able to recover the components because there are more than three eigenvalues. However, if δ is at least 2, reconstruction is guaranteed to work.

We choose the prime ideals $\mathfrak{p}_j = \langle f_j \rangle$, $0 \leq j \leq 2$, generated by the polynomials

$$\begin{aligned} f_0 &= x_1^2 + x_2^2 - 1, \\ f_1 &= (x_1 + 3)^2 + (x_2 - \frac{3}{2})^2 - 1, \\ f_2 &= (x_1 - 2)^2 + (x_2 + 1)^2 - 1. \end{aligned}$$

As $f_0 f_1 \in \mathfrak{p}_0 \cap \mathfrak{p}_1$, but $f_0 f_1 \notin \mathfrak{p}_2$, we infer that we can choose $d = 4$ to meet the requirement $(\mathfrak{p}_k \cap \mathfrak{p}_l \cap R_{\leq d}) \setminus \mathfrak{p}_j \neq \emptyset$ for all $k, l, j \in \{0, 1, 2\}$ with $k \neq l \neq j \neq k$. We set $s := 3$ and choose, for illustration, the weights

$$(\lambda_{ij})_{0 \leq i \leq s, 0 \leq j \leq 2} = \begin{pmatrix} 1 & 1 & 1 \\ 2 & 3 & 6 \\ -1 & 3 & 5 \\ 2 & 1 & -2 \end{pmatrix}, \quad (4.24)$$

which is a matrix of rank 3. The moments of the uniform measure supported on the three circles are rational numbers, up to a normalization constant, and can be computed symbolically as well.

Applying Algorithm 4.3 with $\delta = 0$, we obtain the four eigenvalues

$$[\frac{1}{2} : 1 : -\frac{1}{2} : 1], \quad [1 : 3 : 3 : 1], \quad [-\frac{1}{2} : -3 : -\frac{5}{2} : 1], \quad [3 : 11 : 7 : 1],$$

indeed more than the number of components. The first three of these eigenvalues are homogeneous coordinates for the columns of (4.24), while the fourth eigenvalue has coordinates equal to the sum of the columns of (4.24), which we expected in view of Example 4.3.19. So far, however, we did not know whether this would be the only such extraneous eigenvalue, in this case. The eigenvectors corresponding to the first three eigenvalues are the polynomials $f_1 f_2$, $f_0 f_2$, $f_0 f_1$, respectively, of total degree 4, which are Lagrange-like polynomials in the sense of Remark 4.3.3. The eigenvector for the fourth eigenvalue is the polynomial

$$x_1^4 + 8x_1^3 x_2 + 24x_1^2 x_2^2 + 32x_1 x_2^3 + 16x_2^4 - 5x_1^2 - 20x_1 x_2 - 20x_2^2 + \frac{25}{8},$$

which is the polynomial of interest in Proposition 4.3.18.

Next, applying the algorithm for $\delta = 1$ allows to eliminate the fourth eigenvalue. This means that the sufficient bound on δ from Corollary 4.5.6 (i. e. $\delta \geq 2$ as $\mathfrak{p}_0, \mathfrak{p}_1, \mathfrak{p}_2$ are

generated by polynomials of degree 2) is not sharp in this case. Solving the linear system as described in Remark 4.7.7, we obtain matrices $\tilde{H}_0, \tilde{H}_1, \tilde{H}_2$ (which, up to permutation and scaling, agree with the moment matrices H_0, H_1, H_2 from Example 4.3.19) in terms of the basis consisting of the 21 monomials of total degree at most $d+\delta = 5$. For instance, the kernel of \tilde{H}_0 is spanned by the polynomials

$$\begin{array}{lll} -x_1^4 - 2x_1^2x_2^2 - x_2^4 + 1, & -x_1^4 - x_1^2x_2^2 + x_1^2, & -x_1^5 - x_1^3x_2^2 + x_1^3, \\ -x_1^5 - 2x_1^3x_2^2 - x_1x_2^4 + x_1, & -x_1^3x_2 - x_1x_2^3 + x_1x_2, & -x_1^4x_2 - x_1^2x_2^3 + x_1^2x_2, \\ -x_1^4x_2 - 2x_1^2x_2^3 - x_2^5 + x_2, & -x_1^2x_2^2 - x_2^4 + x_2^2, & -x_1^3x_2^2 - x_1x_2^4 + x_1x_2^2, \\ & & -x_1^2x_2^3 - x_2^5 + x_2^3. \end{array}$$

Using algebraic methods such as the computation of a Gröbner basis, we can determine that these polynomials generate the ideal $\mathfrak{p}_0 = \langle x_1^2 + x_2^2 - 1 \rangle$. \diamond

List of figures

2.1	Piecewise-linear function of Example 2.4.2 with Fourier partial sum approximation	39
2.2	2-mixture of local Gaussian distributions of Example 2.4.6 and its reconstruction	42
3.1	Functions $P_{d,1}$, P_d and $Q_{d,\varepsilon}$ on \mathbf{T}^2	73
3.2	Functions $P_{d,1}$, P_d and $Q_{d,\varepsilon}$ on \mathbf{T}^1	74
4.1	Point sets $V(\mathbf{a}_1)$, $V(\mathbf{a}_2)$ of Example 4.6.1	119
4.2	Reconstruction of 3 trigonometric curves	137
4.3	Weighted sums of measures supported on 6 trigonometric curves	138
4.4	Reconstruction of 6 trigonometric curves	139

The figures have been created with SageMath, Matplotlib and TikZ.

List of algorithms

1.1	Prony's method	13
2.1	Parameter recovery for mixtures of local Dirac mixtures with r components of order l	34
4.1	Recovery of components from several weighted sums	95
4.2	Positive-semidefinite recovery from several weighted sums	110
4.3	Recovery of weights from several weighted sums (degree of generators known)	124
4.4	Recovery of weights from several weighted sums (degree of generators unknown)	128

Glossary of symbols

\mathbb{N}	natural numbers $\{0, 1, 2, \dots\}$
\mathbb{Z}	integers
\mathbb{R}	real numbers
\mathbb{C}	complex numbers
\mathbb{k}	field
\mathbb{k}^*	non-zero elements of \mathbb{k} , algebraic torus
$\mathbb{P}_{\mathbb{k}}^n$	projective space of dimension n
\mathbf{T}	complex torus: unit circle in \mathbb{C}
\mathbf{T}	complex torus: periodic unit interval $[0, 1)$
δ_{ij}	Kronecker delta function of $i, j \in \mathbb{Z}$
$-!$	factorial
$\lfloor - \rfloor$	floor function
$\lceil - \rceil$	ceiling function
$\#-$	cardinality of a set
$\text{char}(-)$	characteristic of a field
$\text{conv}(-)$	convex hull
$\text{diag}(-)$	diagonal matrix
$\text{dim}(-)$	dimension of a vector space or variety
$\text{im}(-)$	image of a map or matrix
$\text{ker}(-)$	kernel of a map or matrix
$\text{rk}(-)$	rank of a matrix or tensor
$\text{supp}(-)$	support of a function, measure or distribution
$\langle - \rangle$	vector subspace or ideal spanned by some elements
π_V	projection map to a vector subspace V
V^\perp	orthogonal complement of a vector subspace V
V^*	algebraic dual space of a vector space V
$\text{Hom}_{\mathbb{k}}(V, \mathbb{k})$	linear maps from a vector space V to \mathbb{k}
$\text{Hom}_{\mathbb{k}}^{\text{semi}}(V, \mathbb{k})$	semilinear maps from a vector space V to \mathbb{k}
\mathbf{I}_n	identity matrix of size $n \times n$
$-^\top$	transpose of a matrix or map
$-^*$	conjugate transpose of a matrix
$\bar{-}$	conjugate of a matrix or vector, residue class
$-^\dagger$	Moore–Penrose pseudo-inverse of a matrix
$ - $	absolute value

$ \alpha $	total degree of $\alpha \in \mathbb{N}^n$
$ \alpha _\infty$	max-degree of $\alpha \in \mathbb{Z}^n$
$ \alpha _R$	degree with respect to a filtration of R
$\ -\ _2$	2-norm of a vector, spectral norm of a matrix
$\ -\ _1$	L^1 -norm of a function
$\ -\ _\infty$	supremum norm, L^∞ -norm of a function
$\ -\ _F$	Frobenius norm of a matrix
$\langle -, - \rangle$	Euclidean scalar product
$f^{(k)}$	k -th derivative of a function f
$C^k(\Omega)$	k -times continuously differentiable functions from Ω to \mathbb{k}
$C_c^0(\Omega)$	compactly-supported continuous functions on Ω
$C_0^0(\Omega)$	continuous functions on Ω vanishing at infinity
$O(-)$	big O
$\Theta(-)$	big Theta
δ_ξ	Dirac measure located at a point $\xi \in \mathbb{k}^n$
ev_ξ	evaluation homomorphism at a point $\xi \in \mathbb{k}^n$
\mathfrak{m}_ξ	ideal associated to a point $\xi \in \mathbb{k}^n$
$\sqrt{-}$	radical of an ideal
$I(-)$	vanishing ideal of a set
$V(-)$	algebraic variety generated by a set of elements
$-\dashrightarrow$	rational map
ν_d	d -uple Veronese embedding
$S^d(-)$	symmetric tensors of order d
$[v_0 : \cdots : v_n]$	homogeneous/projective coordinates of $v \in \mathbb{P}_{\mathbb{k}}^n$
$R_{\leq d}$	d -th component of a filtration of R
$-^\circ$	involution
$\langle -, - \rangle_\sigma$	sesquilinear form associated to a functional σ
$\Phi_{\mathfrak{a}, \sigma}^U, \Phi_{\mathfrak{a}}^U$	sesquilinear map associated to a functional σ , subspace U and ideal \mathfrak{a}
ι	embedding of \mathbf{T}^n into \mathbb{T}^n
$e_{d\xi}$	see Definition 3.5.2
B_d	basis of d -th component of a filtration
F_d	Fejér kernel of degree d
$P_d, P_{d,1}$	see Definition 3.5.4
$Q_{d,0}, Q_{d,\varepsilon}$	see Definition 3.5.4, regularized Christoffel function
Q_d	see Remark 3.5.5, Christoffel function
$\hat{-}$	Fourier transform/coefficients of a function or measure
$\check{-}$	reflection of a function
$- \circ -$	composition of functions
$- * -$	convolution
$- _U$	restriction to a set U
$\mathbb{1}_U$	indicator function of a set U
$\xrightarrow{\text{weak}^*}$	weak* convergence

Bibliography

- [ACdH10] F. Andersson, M. Carlsson, and M. V. de Hoop. “Nonlinear approximation of functions in two dimensions by sums of exponential functions”. In: *Appl. Comput. Harmon. Anal.* 29.2 (2010), pp. 156–181. DOI: [10.1016/j.acha.2009.08.009](https://doi.org/10.1016/j.acha.2009.08.009).
- [AM07] K. Anaya-Izquierdo and P. Marriott. “Local mixture models of exponential families”. In: *Bernoulli* 13.3 (2007), pp. 623–640. DOI: [10.3150/07-BEJ6170](https://doi.org/10.3150/07-BEJ6170).
- [And+99] E. Anderson et al. *LAPACK Users’ Guide*. 3rd ed. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1999. DOI: [10.1137/1.9780898719604](https://doi.org/10.1137/1.9780898719604).
- [Arn83] B. Arnold. *Pareto distributions*. Statistical distributions in scientific work series. International Co-operative Publishing House, 1983.
- [ARS18] C. Améndola, K. Ranestad, and B. Sturmfels. “Algebraic identifiability of Gaussian mixtures”. In: *Int. Math. Res. Not.* 21 (2018), pp. 6556–6580. DOI: [10.1093/imrn/rnx090](https://doi.org/10.1093/imrn/rnx090).
- [Atk72] F. V. Atkinson. *Multiparameter eigenvalue problems*. Volume I: Matrices and compact operators, Mathematics in Science and Engineering, Vol. 82. New York, London: Academic Press, 1972, pp. xii+209.
- [BBCM13] A. Bernardi, J. Brachat, P. Comon, and B. Mourrain. “General tensor decomposition, moment matrices and applications”. In: *J. Symbolic Comput.* 52 (2013), pp. 51–71. DOI: [10.1016/j.jsc.2012.05.012](https://doi.org/10.1016/j.jsc.2012.05.012).
- [BBEM90] B. Beauzamy, E. Bombieri, P. Enflo, and H. L. Montgomery. “Products of polynomials in many variables”. In: *J. Number Theory* 36.2 (1990), pp. 219–245. DOI: [10.1016/0022-314X\(90\)90075-3](https://doi.org/10.1016/0022-314X(90)90075-3).
- [BC13] P. Bürgisser and F. Cucker. *Condition. The geometry of numerical algorithms*. Vol. 349. Grundlehren math. Wissensch. Heidelberg: Springer, 2013, pp. xxxii+554. DOI: [10.1007/978-3-642-38896-5](https://doi.org/10.1007/978-3-642-38896-5).
- [BCGI07] A. Bernardi, M. V. Catalisano, A. Gimigliano, and M. Idà. “Osculating varieties of Veronese varieties and their higher secant varieties”. In: *Canad. J. Math.* 59.3 (2007), pp. 488–502. DOI: [10.4153/CJM-2007-021-6](https://doi.org/10.4153/CJM-2007-021-6).
- [BCMT10] J. Brachat, P. Comon, B. Mourrain, and E. Tsigaridas. “Symmetric tensor decomposition”. In: *Linear Algebra Appl.* 433.11 (2010), pp. 1851–1872. DOI: [10.1016/j.laa.2010.06.046](https://doi.org/10.1016/j.laa.2010.06.046).
- [BDGY21] D. Batenkov, B. Diederichs, G. Goldman, and Y. Yomdin. “The spectral properties of Vandermonde matrices with clustered nodes”. In: *Linear Algebra Appl.* 609 (2021), pp. 37–72. DOI: [10.1016/j.laa.2020.08.034](https://doi.org/10.1016/j.laa.2020.08.034).

BIBLIOGRAPHY

- [BEKS17] J. Bezanson, A. Edelman, S. Karpinski, and V. Shah. “Julia: A Fresh Approach to Numerical Computing”. In: *SIAM Review* 59.1 (2017), pp. 65–98. DOI: [10.1137/141000671](https://doi.org/10.1137/141000671).
- [Bjö96] Å. Björck. *Numerical methods for least squares problems*. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM), 1996, pp. xviii+408. DOI: [10.1137/1.9781611971484](https://doi.org/10.1137/1.9781611971484).
- [Bou06] N. Bourbaki. *Algèbre commutative. Chapitres 1 à 4*. 2nd ed. Éléments de mathématique. Berlin, Heidelberg: Springer, 2006, pp. vi+356. DOI: [10.1007/978-3-540-33976-2](https://doi.org/10.1007/978-3-540-33976-2).
- [BT18] P. Breiding and S. Timme. “HomotopyContinuation.jl: A Package for Homotopy Continuation in Julia”. In: *Mathematical Software – ICMS 2018*. Cham: Springer International Publishing, 2018, pp. 458–465. DOI: [10.1007/978-3-319-96418-8_54](https://doi.org/10.1007/978-3-319-96418-8_54).
- [BT20] A. Bernardi and D. Taufer. “Waring, tangential and cactus decompositions”. In: *J. Math. Pures Appl.* 9th ser. 143 (2020), pp. 1–30. DOI: [10.1016/j.matpur.2020.07.003](https://doi.org/10.1016/j.matpur.2020.07.003).
- [Car21] R. D. Carmichael. “Boundary Value and Expansion Problems: Algebraic Basis of the Theory”. In: *Amer. J. Math.* 43.2 (1921), pp. 69–101. DOI: [10.2307/2370243](https://doi.org/10.2307/2370243).
- [CGG02] M. V. Catalisano, A. V. Geramita, and A. Gimigliano. “On the secant varieties to the tangential varieties of a Veronesean”. In: *Proc. Amer. Math. Soc.* 130.4 (2002), pp. 975–985. DOI: [10.1090/S0002-9939-01-06251-7](https://doi.org/10.1090/S0002-9939-01-06251-7).
- [CLO15] D. A. Cox, J. Little, and D. O’Shea. *Ideals, varieties, and algorithms*. 4th ed. Undergraduate Texts in Mathematics. An introduction to computational algebraic geometry and commutative algebra. Cham: Springer, 2015, pp. xvi+646. DOI: [10.1007/978-3-319-16721-3](https://doi.org/10.1007/978-3-319-16721-3).
- [Con90] J. B. Conway. *A course in functional analysis*. 2nd ed. Vol. 96. Graduate Texts in Mathematics. Corr. fourth print. New York: Springer, 1990, pp. xvi+399.
- [COV17] L. Chiantini, G. Ottaviani, and N. Vannieuwenhoven. “On generic identifiability of symmetric tensors of subgeneric rank”. In: *Trans. Amer. Math. Soc.* 369.6 (2017), pp. 4021–4042. DOI: [10.1090/tran/6762](https://doi.org/10.1090/tran/6762).
- [CPS12] V. Cortés, J. M. Peña, and T. Sauer. “Simultaneous triangularization of commuting matrices for the solution of polynomial equations”. In: *Cent. Eur. J. Math.* 10.1 (2012), pp. 277–291. DOI: [10.2478/s11533-011-0106-z](https://doi.org/10.2478/s11533-011-0106-z).
- [Dem00] J. Demmel. “A Brief Tour of Eigenproblems”. In: *Templates for the Solution of Algebraic Eigenvalue Problems*. 2000. Chap. 2, pp. 7–36. DOI: [10.1137/1.9780898719581.ch2](https://doi.org/10.1137/1.9780898719581.ch2).
- [Eis92] D. Eisenbud. “Green’s conjecture: an orientation for algebraists”. In: *Free resolutions in commutative algebra and algebraic geometry (Sundance, UT, 1990)*. Vol. 2. Res. Notes Math. Boston, MA: Jones and Bartlett, 1992, pp. 51–78.

BIBLIOGRAPHY

- [Eis99] D. Eisenbud. *Commutative algebra with a view toward algebraic geometry*. Corr. third print. Vol. 150. Graduate Texts in Mathematics. New York: Springer, 1999, pp. xvi+797. DOI: [10.1007/978-1-4612-5350-1](https://doi.org/10.1007/978-1-4612-5350-1).
- [FAV16] M. Fatemi, A. Amini, and M. Vetterli. “Sampling and reconstruction of shapes with algebraic boundaries”. In: *IEEE Trans. Signal Process.* 64.22 (2016), pp. 5807–5818. DOI: [10.1109/TSP.2016.2591505](https://doi.org/10.1109/TSP.2016.2591505).
- [Gra14] L. Grafakos. *Classical Fourier analysis*. 3rd ed. Vol. 249. Graduate Texts in Mathematics. New York: Springer, 2014, pp. xviii+638. DOI: [10.1007/978-1-4939-1194-3](https://doi.org/10.1007/978-1-4939-1194-3).
- [Gro21] A. Grosdos Koutsoumpelias. “Algebraic Methods for the Estimation of Statistical Distributions”. PhD thesis. Osnabrück University, 2021. URL: <https://nbn-resolving.org/urn:nbn:de:gbv:700-202107155198>.
- [GV96] G. H. Golub and C. F. Van Loan. *Matrix computations*. 3rd ed. Johns Hopkins Studies in the Mathematical Sciences. Baltimore, MD: Johns Hopkins University Press, 1996, pp. xxx+698.
- [GW20] A. Grosdos Koutsoumpelias and M. Wageringel. “Moment ideals of local Dirac mixtures”. In: *SIAM J. Appl. Algebra Geom.* 4.1 (2020), pp. 1–27. DOI: [10.1137/18M1219783](https://doi.org/10.1137/18M1219783).
- [Hac19] W. Hackbusch. *Tensor spaces and numerical tensor calculus*. 2nd ed. Vol. 56. Springer Series in Computational Mathematics. Cham: Springer, 2019, pp. xxviii+605. DOI: [10.1007/978-3-030-35554-8](https://doi.org/10.1007/978-3-030-35554-8).
- [Har+20] C. R. Harris et al. “Array programming with NumPy”. In: *Nature* 585.7825 (2020), pp. 357–362. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2).
- [Har87] R. Hartshorne. *Algebraic geometry*. Fourth printing. Vol. 52. Graduate Texts in Mathematics. New York, Heidelberg: Springer, 1987, pp. xvi+496.
- [HJ13] R. A. Horn and C. R. Johnson. *Matrix analysis*. 2nd ed. Cambridge: Cambridge University Press, 2013, pp. xviii+643.
- [HJ94] R. A. Horn and C. R. Johnson. *Topics in matrix analysis*. Corrected reprint of the 1991 original. Cambridge: Cambridge University Press, 1994, pp. viii+607.
- [HKP04] M. E. Hochstenbach, T. Košir, and B. Plestenjak. “A Jacobi–Davidson type method for the two-parameter eigenvalue problem”. In: *SIAM J. Matrix Anal. Appl.* 26.2 (2004), pp. 477–497. DOI: [10.1137/S0895479802418318](https://doi.org/10.1137/S0895479802418318).
- [HS90] Y. Hua and T. K. Sarkar. “Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise”. In: *IEEE Trans. Acoust. Speech Signal Process.* 38.5 (1990), pp. 814–824. DOI: [10.1109/29.56027](https://doi.org/10.1109/29.56027).
- [IK99] A. Iarrobino and V. Kanev. *Power Sums, Gorenstein Algebras, and Determinantal Loci*. Lecture Notes in Mathematics. Berlin, Heidelberg: Springer, 1999.
- [Isk18] A. Iske. *Approximation theory and algorithms for data analysis*. Vol. 68. Texts in Applied Mathematics. Original German edition published by

- Springer-Verlag, Heidelberg, 2017. Cham: Springer, 2018, pp. x+358. DOI: [10.1007/978-3-030-05228-7](https://doi.org/10.1007/978-3-030-05228-7).
- [Joh17] F. Johansson. “Arb: efficient arbitrary-precision midpoint-radius interval arithmetic”. In: *IEEE Transactions on Computers* 66.8 (2017), pp. 1281–1292. DOI: [10.1109/TC.2017.2690633](https://doi.org/10.1109/TC.2017.2690633).
- [Kad18] V. Kadets. *A course in functional analysis and measure theory*. Universitext. Cham: Springer, 2018, pp. xxii+539. DOI: [10.1007/978-3-319-92004-7](https://doi.org/10.1007/978-3-319-92004-7).
- [Käm13] L. Kämmerer. “Reconstructing hyperbolic cross trigonometric polynomials by sampling along rank-1 lattices”. In: *SIAM J. Numer. Anal.* 51.5 (2013), pp. 2773–2796. DOI: [10.1137/120871183](https://doi.org/10.1137/120871183).
- [KMvdO19] S. Kunis, H. M. Möller, and U. von der Ohe. “Prony’s method on the sphere”. In: *SMAI J. Comput. Math.* S5 (2019), pp. 87–97. DOI: [10.5802/smai-jcm.53](https://doi.org/10.5802/smai-jcm.53).
- [KN20] S. Kunis and D. Nagel. “On the smallest singular value of multivariate Vandermonde matrices with clustered nodes”. In: *Linear Algebra Appl.* 604 (2020), pp. 1–20. DOI: [10.1016/j.laa.2020.06.003](https://doi.org/10.1016/j.laa.2020.06.003).
- [KPRvdO16] S. Kunis, T. Peter, T. Römer, and U. von der Ohe. “A multivariate generalization of Prony’s method”. In: *Linear Algebra Appl.* 490 (2016), pp. 31–47. DOI: [10.1016/j.laa.2015.10.023](https://doi.org/10.1016/j.laa.2015.10.023).
- [KRvdO20] S. Kunis, T. Römer, and U. von der Ohe. “Learning algebraic decompositions using Prony structures”. In: *Adv. Appl. Math.* 118 (2020). Article 102044, pp. 1–43. DOI: [10.1016/j.aam.2020.102044](https://doi.org/10.1016/j.aam.2020.102044).
- [KV96] H. Krim and M. Viberg. “Two decades of array signal processing research: the parametric approach”. In: *IEEE Signal Processing Magazine* 13.4 (1996), pp. 67–94. DOI: [10.1109/79.526899](https://doi.org/10.1109/79.526899).
- [Lan12] J. M. Landsberg. *Tensors: geometry and applications*. Vol. 128. Graduate Studies in Mathematics. Providence, RI: American Mathematical Society, 2012, pp. xx+439.
- [Li03] T. Y. Li. “Numerical solution of polynomial systems by homotopy continuation methods”. In: *Handbook of numerical analysis, Vol. XI*. Handb. Numer. Anal., XI. Amsterdam: North-Holland, 2003, pp. 209–304.
- [LP15] J. B. Lasserre and M. Putinar. “Algebraic-exponential data recovery from moments”. In: *Discrete Comput. Geom.* 54.4 (2015), pp. 993–1012. DOI: [10.1007/s00454-015-9739-1](https://doi.org/10.1007/s00454-015-9739-1).
- [LP19] J. B. Lasserre and E. Pauwels. “The empirical Christoffel function with applications in data analysis”. In: *Adv. Comput. Math.* 45.3 (2019), pp. 1439–1468. DOI: [10.1007/s10444-019-09673-1](https://doi.org/10.1007/s10444-019-09673-1).
- [LR12] M. Laurent and P. Rostalski. “The approach of moments for polynomial equations”. In: *Handbook on semidefinite, conic and polynomial optimization*. Vol. 166. Internat. Ser. Oper. Res. Management Sci. New York: Springer, 2012, pp. 25–60. DOI: [10.1007/978-1-4614-0769-0_2](https://doi.org/10.1007/978-1-4614-0769-0_2).

- [LRA93] S. E. Leurgans, R. T. Ross, and R. B. Abel. “A decomposition for three-way arrays”. In: *SIAM J. Matrix Anal. Appl.* 14.4 (1993), pp. 1064–1083. DOI: [10.1137/0614071](https://doi.org/10.1137/0614071).
- [LZGL21] W. Li, Z. Zhu, W. Gao, and W. Liao. *Stability and Super-resolution of MUSIC and ESPRIT for Multi-snapshot Spectral Estimation*. 2021. arXiv: [2105.14304](https://arxiv.org/abs/2105.14304) [cs.IT].
- [Mac95] I. G. Macdonald. *Symmetric functions and Hall polynomials*. 2nd ed. Oxford Mathematical Monographs. With contributions by A. Zelevinsky, Oxford Science Publications. New York: The Clarendon Press, Oxford University Press, 1995, pp. x+475.
- [Mar+21] S. Marx, E. Pauwels, T. Weisser, D. Henrion, and J. B. Lasserre. “Semi-algebraic Approximation Using Christoffel–Darboux Kernel”. In: *Constructive Approximation* (2021). DOI: [10.1007/s00365-021-09535-4](https://doi.org/10.1007/s00365-021-09535-4).
- [Mar02] P. Marriott. “On the local geometry of mixture models”. In: *Biometrika* 89 (2002), pp. 77–93.
- [Mey00] C. Meyer. *Matrix analysis and applied linear algebra*. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM), 2000, pp. xii+718. DOI: [10.1137/1.9780898719512](https://doi.org/10.1137/1.9780898719512).
- [MNT91] A. Máté, P. Nevai, and V. Totik. “Szegő’s extremum problem on the unit circle”. In: *Ann. Math.* 2nd ser. 134.2 (1991), pp. 433–453. DOI: [10.2307/2944352](https://doi.org/10.2307/2944352).
- [Moi15] A. Moitra. “Super-resolution, extremal functions and the condition number of Vandermonde matrices”. In: *STOC’15—Proceedings of the 2015 ACM Symposium on Theory of Computing*. New York: ACM, 2015, pp. 821–830. DOI: [10.1145/2746539.2746561](https://doi.org/10.1145/2746539.2746561).
- [Mou18] B. Mourrain. “Polynomial–Exponential Decomposition From Moments”. In: *Found. Comput. Math.* 18.6 (2018), pp. 1435–1492. DOI: [10.1007/s10208-017-9372-x](https://doi.org/10.1007/s10208-017-9372-x).
- [MP09] A. Muhič and B. Plestenjak. “On the singular two-parameter eigenvalue problem”. In: *Electron. J. Linear Algebra* 18 (2009), pp. 420–437. DOI: [10.13001/1081-3810.1322](https://doi.org/10.13001/1081-3810.1322).
- [MSW21] H. Melánová, B. Sturmfels, and R. Winter. *Recovery from Power Sums*. 2021. arXiv: [2106.13981](https://arxiv.org/abs/2106.13981) [math.AG].
- [Nev86] P. Nevai. “Géza Freud, orthogonal polynomials and Christoffel functions. A case study”. In: *J. Approx. Theory* 48.1 (1986), pp. 3–167. DOI: [10.1016/0021-9045\(86\)90016-X](https://doi.org/10.1016/0021-9045(86)90016-X).
- [OJ15] G. Ongie and M. Jacob. “Recovery of piecewise smooth images from few Fourier samples”. In: *2015 International Conference on Sampling Theory and Applications (SampTA)*. 2015, pp. 543–547. DOI: [10.1109/SAMP TA.2015.7148950](https://doi.org/10.1109/SAMP TA.2015.7148950).
- [OJ16] G. Ongie and M. Jacob. “Off-the-grid recovery of piecewise constant images from few Fourier samples”. In: *SIAM J. Imaging Sci.* 9.3 (2016), pp. 1004–1041. DOI: [10.1137/15M1042280](https://doi.org/10.1137/15M1042280).

- [PPL21] E. Pauwels, M. Putinar, and J.-B. Lasserre. “Data analysis from empirical moments and the Christoffel function”. In: *Found. Comput. Math.* 21.1 (2021), pp. 243–273. DOI: [10.1007/s10208-020-09451-2](https://doi.org/10.1007/s10208-020-09451-2).
- [PPST18] G. Plonka, D. Potts, G. Steidl, and M. Tasche. *Numerical Fourier analysis*. Applied and Numerical Harmonic Analysis. Cham: Birkhäuser/Springer, 2018, pp. xvi+168. DOI: [10.1007/978-3-030-04306-3](https://doi.org/10.1007/978-3-030-04306-3).
- [Pro95] R. Prony. “Essai expérimental et analytique: Sur les lois de la Dilatabilité de fluides élastique et sur celles de la Force expansive de la vapeur de l’eau et de la vapeur de l’alkool, à différentes températures”. In: *Journal de l’école polytechnique* 2 (1795), pp. 24–76. URL: <https://gallica.bnf.fr/ark:/12148/bpt6k433661n>.
- [PT14] G. Plonka and M. Tasche. “Prony methods for recovery of structured functions”. In: *GAMM-Mitteilungen* 37.2 (2014), pp. 239–258. DOI: [10.1002/gamm.201410011](https://doi.org/10.1002/gamm.201410011).
- [Rin17] C. M. Ringel. “The eigenvector variety of a matrix pencil”. In: *Linear Algebra Appl.* 531 (2017), pp. 447–458. DOI: [10.1016/j.laa.2017.05.004](https://doi.org/10.1016/j.laa.2017.05.004).
- [RK90] R. Roy and T. Kailath. “ESPRIT—estimation of signal parameters via rotational invariance techniques”. In: *Signal Processing, Part II*. Vol. 23. IMA Vol. Math. Appl. New York: Springer, 1990, pp. 369–411.
- [Rud87] W. Rudin. *Real and complex analysis*. 3rd ed. New York: McGraw-Hill Book Co., 1987, pp. xiv+416.
- [Sag21] The Sage Developers. *SageMath, the Sage Mathematics Software System*. Version 9.4. 2021. URL: <https://www.sagemath.org>.
- [Sau17] T. Sauer. “Prony’s method in several variables”. In: *Numer. Math.* 136.2 (2017), pp. 411–438. DOI: [10.1007/s00211-016-0844-8](https://doi.org/10.1007/s00211-016-0844-8).
- [Sau18] T. Sauer. “Prony’s method in several variables: symbolic solutions by universal interpolation”. In: *J. Symbolic Comput.* 84 (2018), pp. 95–112. DOI: [10.1016/j.jsc.2017.03.006](https://doi.org/10.1016/j.jsc.2017.03.006).
- [Sch17] K. Schmüdgen. *The moment problem*. Vol. 277. Graduate Texts in Mathematics. Cham: Springer, 2017, pp. xii+535.
- [Sch73] L. Schwartz. *Théorie des distributions*. 3. nouv. éd. Paris: Hermann, 1973.
- [Sch86] R. Schmidt. “Multiple emitter location and signal parameter estimation”. In: *IEEE Transactions on Antennas and Propagation* 34.3 (1986), pp. 276–280. DOI: [10.1109/TAP.1986.1143830](https://doi.org/10.1109/TAP.1986.1143830).
- [SG12] J. Skowron and A. Gould. *General Complex Polynomial Root Solver and Its Further Optimization for Binary Microlenses*. 2012. arXiv: [1203.1034](https://arxiv.org/abs/1203.1034) [[astro-ph.EP](https://arxiv.org/abs/1203.1034)].
- [Shi18] Y. Shitov. “A Counterexample to Comon’s Conjecture”. In: *SIAM Journal on Applied Algebra and Geometry* 2.3 (2018), pp. 428–443. DOI: [10.1137/17M1131970](https://doi.org/10.1137/17M1131970).
- [Sle+00] G. Sleijpen, H. van der Vorst, A. Ruhe, Z. Bai, T. Ericsson, T. Kowalski, B. Kågström, and R. Li. “Generalized Non-Hermitian Eigenvalue Prob-

- lems”. In: *Templates for the Solution of Algebraic Eigenvalue Problems*. 2000. Chap. 8, pp. 233–279. DOI: [10.1137/1.9780898719581.ch8](https://doi.org/10.1137/1.9780898719581.ch8).
- [Sy186] J. J. Sylvester. “Sur une extension d’un théorème de Clebsch relatif aux courbes du quatrième degré”. In: *Comptes rendus de l’Académie des Sciences Paris* 102, Jan.–June (1886), pp. 1532–1534. URL: <https://gallica.bnf.fr/ark:/12148/bpt6k3058f>.
- [Tsa+20] M. C. Tsakiris, L. Peng, A. Conca, L. Kneip, Y. Shi, and H. Choi. “An Algebraic-Geometric Approach for Linear Regression Without Correspondences”. In: *IEEE Transactions on Information Theory* 66.8 (2020), pp. 5130–5144. DOI: [10.1109/TIT.2020.2977166](https://doi.org/10.1109/TIT.2020.2977166).
- [Van83] P. Van Dooren. “Reducing subspaces: Definitions, properties and algorithms”. In: *Matrix Pencils*. Ed. by B. Kågström and A. Ruhe. Berlin, Heidelberg: Springer, 1983, pp. 58–73. DOI: [10.1007/BFb0062094](https://doi.org/10.1007/BFb0062094).
- [VB96] L. Vandenberghe and S. Boyd. “Semidefinite programming”. In: *SIAM Rev.* 38.1 (1996), pp. 49–95. DOI: [10.1137/1038003](https://doi.org/10.1137/1038003).
- [vdOhe17] U. von der Ohe. “On the reconstruction of multivariate exponential sums”. PhD thesis. Osnabrück University, 2017. URL: <https://nbn-resolving.org/urn:nbn:de:gbv:700-2017120716391>.
- [vdOhe21] U. von der Ohe. “Toward a structural theory of learning algebraic decompositions”. PhD thesis. Università degli Studi di Genova, 2021. URL: <https://hdl.handle.net/11567/1045060>.
- [Vir+20] P. Virtanen et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17 (2020), pp. 261–272. DOI: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2).
- [Wag21] M. Wageringel. *Momentproblems*. Version 1.0.0. 2021. URL: <https://github.com/mwageringel/momentproblems>.
- [Wol18] Wolfram Research, Inc. *Mathematica 11*. Version 11.3. 2018. URL: <https://www.wolfram.com>.