



The empirical Christoffel function with applications in data analysis

Jean B. Lasserre¹ · Edouard Pauwels² 

Received: 18 December 2017 / Accepted: 4 February 2019 /

Published online: 7 March 2019

© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

We illustrate the potential applications in machine learning of the Christoffel function, or, more precisely, its empirical counterpart associated with a counting measure uniformly supported on a finite set of points. Firstly, we provide a thresholding scheme which allows approximating the support of a measure from a finite subset of its moments with strong asymptotic guaranties. Secondly, we provide a consistency result which relates the empirical Christoffel function and its population counterpart in the limit of large samples. Finally, we illustrate the relevance of our results on simulated and real-world datasets for several applications in statistics and machine learning: (a) density and support estimation from finite samples, (b) outlier and novelty detection, and (c) affine matching.

Keywords Christoffel function · Statistics · Support inference · Density estimation · Consistency

Mathematics Subject Classification (2010) 62-07 · 62H99 · 68T05

1 Introduction

The main claim of this paper is that the Christoffel function (a tool from Approximation Theory) can prove to be very useful in machine learning applications. The Christoffel function is associated with a finite measure and a degree parameter d . It has an important history of research with a strong connection to orthogonal polynomials [16, 40], and interpolation and approximation theory [14, 30]. Its typical

Communicated by: Tomas Sauer

✉ Edouard Pauwels
edouard.pauwels@irit.fr

¹ LAAS-CNRS and Institute of Mathematics, University of Toulouse, LAAS, 7 avenue du Colonel Roche, 31077 Toulouse, France

² IRIT, Université Toulouse 3 Paul Sabatier, 118 route de Narbonne, 31062 Toulouse, France

asymptotic behavior as d increases is of particular interest because, in some specific settings, it provides very relevant information on the support and the density of the associated input measure. Important references include [18, 27, 28, 41] in a single dimension, [5–7, 22, 46] for specific multivariate settings, and [21] for ratios of mutually absolutely continuous measures. The topic is still a subject of active research, but regarding properties of the Christoffel function, a lot of information is already available.

The present work shows how properties of the Christoffel function can be used successfully in some machine learning applications. To the best of our knowledge, this is the first attempt in such a context with the recent work of [23] and [25]. More precisely, we consider the *empirical* Christoffel function, a specific case where the input measure is a scaled counting measure uniformly supported on a set (a cloud) of data-points. This methodology has three distinguishing features: (i) It is extremely simple and involves no optimization procedure; (ii) it scales linearly with the number of observations (one pass over the data is sufficient); and (iii) it is affine invariant. These three features prove to be especially important in all the applications that we consider.

In [23], we have exhibited a striking property of some distinguished family of sum-of-squares (SOS) polynomials $(Q_d)_{d \in \mathbb{N}}$, indexed by their degree ($2d \in \mathbb{N}$), and easily computed from empirical moments associated with a cloud of n points in \mathbb{R}^p which we call \mathbf{X} . The associated family of sublevel sets $S_{\alpha,d} = \{\mathbf{x} : Q_d(\mathbf{x}) \leq \alpha\}$, for various values of $\alpha > 0$, approximates the global shape of original cloud of points \mathbf{X} . The degree index d can be used as a tuning parameter, trading off regularity of the polynomial Q_d with the fitness of the approximation of the shape of \mathbf{X} (as long as the cloud contains sufficiently many points). Remarkably, even with relatively low degree d , the sets $S_{\alpha,d}$ capture accurately the shape of \mathbf{X} , and so provides a compact (algebraic) encoding of the cloud.

In fact, the reciprocal function $\mathbf{x} \mapsto Q_d(\mathbf{x})^{-1}$ is precisely the *Christoffel function* $\Lambda_{\mu_n,d}$ associated to the empirical counting measure supported on \mathbf{X} and the degree index d . Some properties of the Christoffel function stemming from approximation theory suggest that it could be exploited in a statistical learning context by considering its empirical counterpart. The purpose of this work is to push this idea further. In particular, we investigate (a) further properties of the Christoffel function which prove to be relevant in some machine learning applications, (b) statistical properties of the empirical Christoffel function, and (c) further applications to well-known machine learning tasks.

Contributions

This paper significantly extends [23] in several directions. Indeed our contribution is threefold:

- I. We first provide a thresholding scheme which allows approximating the compact support S of a measure with strong asymptotic guarantees. This result rigorously establishes the property that, as d increases, the scaled Christoffel function decreases to 0 outside S and remains positive in the interior of S .

- II. In view of potential applications in machine learning, we provide a rationale for using the empirical Christoffel function in place of its population counterpart in the limit of a large sample size. We consider a compactly supported population measure μ as well as an empirical measure μ_n uniformly supported on a sample of n vectors in \mathbb{R}^p , drawn independently from μ . For each fixed d , we show a highly desirable strong asymptotic property as n increases. Namely, the empirical Christoffel function $\Lambda_{\mu_n, d}(\cdot)$ converges, *uniformly in $\mathbf{x} \in \mathbb{R}^p$* , to $\Lambda_{\mu, d}(\cdot)$, almost surely with respect to the draw of the random sample.
- III. We illustrate the benefits of the empirical Christoffel function in some important applications, mainly in machine learning. The rationale for such benefits builds on approximation properties of the Christoffel function combined with our consistency result. In particular, we first show on simulated data that the Christoffel function can be useful for *density estimation* and *support inference*. In [23], we have described how the Christoffel function yields a simple procedure for intrusion detection in networks, and here we extend these results by performing a numerical comparison with well-established methods for novelty detection on a real-world dataset. Finally, we show that the Christoffel function is also very useful in performing affine matching and inverse affine shuffling of a dataset.

Comparison with existing literature on set estimation

Support estimation, and more generally set estimation, has a long history in statistics and we intend to give a nonexhaustive overview in this section. The main question of interest is that of inferring a set (support, level sets of the density function...) based on independent samples from an unknown distribution. Pioneering works include [17, 34] followed by [8, 15] and resulted in the introduction and first analyses for estimators based on convex hull for convex domains or union of balls for nonconvex sets. This motivated the development of minimax statistical analysis for the set estimation problem [19, 26, 42] and the introduction of more sophisticated optimal estimators, such as the excess mass estimator [33]. Strong relations between set estimation and density estimation lead to the development of the plugin approach for support and density level set estimation [11, 29] with further generalization proposed in [12] and a precise minimax analysis described in [35].

These works provide a rich statistical analysis of the main estimation approaches currently available. The topic is still active with more precise questions ranging from inference of topological properties [1], new geometric conditions [9], and adaptivity to local properties of the underlying density [32, 39].

One of the goals of our work is the introduction of the Christoffel function as a tool to solve similar problems. This approach has several advantages:

- The Christoffel function allows encoding the global shape of a cloud of points in any finite dimension using a polynomial level set. This kind of encoding is relatively simple and compact. This has clear advantages, for example, the evaluation of a polynomial has a complexity which does not depend on the size of the sample used to calibrate its coefficients and the boundary of the corresponding

sublevel set as a very compact representation as an algebraic set. Furthermore, it turns out that the estimation of the empirical Christoffel function has a computational cost which is linear in the sample size. This is in contrast with distance-based approaches for which membership evaluation requires querying all the sample points. As pointed out in [3], the practical use of multidimensional set estimation techniques involves formidable computational difficulties so that simplicity arises as a major advantage in this context.

- The proposed approach is specific in the sense that it relies on tools which were not considered before for support estimation such as orthogonal polynomials. Topological properties of the support of the distribution or its boundary arise as major questions beyond minimax analysis [1]. In this realm, the objects which we manipulate have a simple algebraic description and could be coupled with computational real algebraic geometry tools to infer topological properties such as, for example, Betty numbers [4]. This strong algebraic structure could in principle allow pushing further the statistical settings which could be handled, with, for example, notions such as singular measures and intrinsic dimension.

We see these facts as potential advantages of the Christoffel function in the context of support estimation and relevant motivation to further study the potential of this procedure in modern data analysis contexts. However, we emphasize that this work constitutes only a first step in this direction. Indeed, we are not able to provide a complete statistical efficiency analysis as precisely described in the support and set estimation literature (e.g., [12]). This would require further studies of precise properties of the Christoffel function itself which are not available given the state of knowledge for this object. We aim at providing a rationale for the proposed approach and motivation for future studies, among which a complete statistical analysis is a longer term goal.¹

Organisation of the paper

Section 2 describes the notation and definitions which will be used throughout the paper. In Section 3, we introduce the Christoffel function, outline some of its known properties, and describe our main theoretical results. Applications are presented in Section 4 where we consider both simulated and real-world data as well as a comparison with well-established machine learning methods. For clarity of exposition, most proofs and technical details are postponed to the [Appendix](#) in Section 6.

2 Notation, definitions, and preliminary results

We fix the ambient dimension to be p throughout the text. For example, we will manipulate vectors in \mathbb{R}^p as well as p -variate polynomials with real coefficients.

¹In particular, qualibration of the underlying polynomial degree as a function of the sample size is out of the scope of this paper and left for future research.

We denote by X a set of p variables X_1, \dots, X_p which we will use in mathematical expressions defining polynomials. We identify monomials from the canonical basis of p -variate polynomials with their exponents in \mathbb{N}^p : we associate to $\alpha = (\alpha_i)_{i=1\dots p} \in \mathbb{N}^p$ the monomial $X^\alpha := X_1^{\alpha_1} X_2^{\alpha_2} \dots X_p^{\alpha_p}$ which degree is $\text{deg}(\alpha) := \sum_{i=1}^p \alpha_i = |\alpha|$. We use the expressions \prec_{gl} and \leq_{gl} to denote the graded lexicographic order, a well ordering over p -variate monomials. This amounts to, first, use the canonical order on the degree and, second, break ties in monomials with the same degree using the lexicographic order with $X_1 = a, X_2 = b \dots$. For example, the monomials in two variables X_1, X_2 , of degree less or equal to 3 listed in this order are given by $1, X_1, X_2, X_1^2, X_1X_2, X_2^2, X_1^3, X_1^2X_2, X_1X_2^2, X_2^3$.

We denote by \mathbb{N}_d^p , the set $\{\alpha \in \mathbb{N}^p; \text{deg}(\alpha) \leq d\}$ ordered by \leq_{gl} . $\mathbb{R}[X]$ denotes the set of p -variate polynomials: linear combinations of monomials with real coefficients. The degree of a polynomial is the highest of the degrees of its monomials with nonzero coefficients.² We use the same notation, $\text{deg}(\cdot)$, to denote the degree of a polynomial or of an element of \mathbb{N}^p . For $d \in \mathbb{N}$, $\mathbb{R}[X]_d$ denotes the set of p -variate polynomials of degree at most d . We set $s(d) = \binom{p+d}{d}$, the number of monomials of degree less or equal to d .

We will denote by $\mathbf{v}_d(X)$ the vector of monomials of degree less or equal to d sorted by \leq_{gl} , i.e., $\mathbf{v}_d(X) := (X^\alpha)_{\alpha \in \mathbb{N}_d^p} \in \mathbb{R}[X]_d^{s(d)}$. With this notation, we can write a polynomial $P \in \mathbb{R}[X]_d$ as $P(X) = \langle \mathbf{p}, \mathbf{v}_d(X) \rangle$ for some real vector of coefficients $\mathbf{p} = (p_\alpha)_{\alpha \in \mathbb{N}_d^p} \in \mathbb{R}^{s(d)}$ ordered using \leq_{gl} . Given $\mathbf{x} = (x_i)_{i=1\dots p} \in \mathbb{R}^p$, $P(\mathbf{x})$ denotes the evaluation of P with the assignments $X_1 = x_1, X_2 = x_2, \dots, X_p = x_p$. Given a Borel probability measure μ and $\alpha \in \mathbb{N}^p$, $y_\alpha(\mu)$ denotes the moment α of μ , i.e., $y_\alpha(\mu) = \int_{\mathbb{R}^p} \mathbf{x}^\alpha d\mu(\mathbf{x})$. Finally, for $\delta > 0$ and every $\mathbf{x} \in \mathbb{R}^p$, let $\mathbf{B}_\delta(\mathbf{x}) := \{\mathbf{x} : \|\mathbf{x}\| \leq \delta\}$ be the closed Euclidean ball of radius δ and centered at \mathbf{x} . We use the shorthand notation \mathbf{B} to denote the closed Euclidean unit ball. For a given subset of Euclidean space, A , ∂A denotes the topological boundary of A . Recall that its Lebesgue volume $\text{vol}(\mathbf{B}_\delta(\mathbf{x}))$ satisfies:

$$\text{vol}(\mathbf{B}_\delta(\mathbf{x})) = \frac{\pi^{\frac{p}{2}}}{\Gamma(\frac{p}{2} + 1)} \delta^p, \quad \forall \mathbf{x} \in \mathbb{R}^p.$$

Furthermore, let $\omega_p := \frac{2\pi^{\frac{p+1}{2}}}{\Gamma(\frac{p+1}{2})}$ denote the surface of the p dimensional unit sphere in \mathbb{R}^{p+1} . Throughout the paper, we will only consider measures of which all moments are finite.

Moment matrix

For a finite Borel measure μ on \mathbb{R}^p denote by $\text{supp}(\mu)$ its support, i.e., the smallest closed set $\Omega \subset \mathbb{R}^p$ such that $\mu(\mathbb{R}^p \setminus \Omega) = 0$. The moment matrix of μ , $\mathbf{M}_d(\mu)$, is a matrix indexed by monomials of degree at most d ordered by \leq_{gl} . For $\alpha, \beta \in \mathbb{N}_d^p$,

²For the null polynomial, we use the convention that its degree is 0 and it is \leq_{gl} smaller than all other monomials.

the corresponding entry in $\mathbf{M}_d(\mu)$ is defined by $\mathbf{M}_d(\mu)_{\alpha,\beta} := y_{\alpha+\beta}(\mu)$, the moment $\int \mathbf{x}^{\alpha+\beta} d\mu$ of μ . When $p = 2$ and $d = 2$, letting $y_\alpha = y_\alpha(\mu)$ for $\alpha \in \mathbb{N}_4^2$, we have:

$$\mathbf{M}_2(\mu) : \begin{matrix} & 1 & X_1 & X_2 & X_1^2 & X_1X_2 & X_2^2 \\ \begin{matrix} 1 \\ X_1 \\ X_2 \\ X_1^2 \\ X_1X_2 \\ X_2^2 \end{matrix} & \begin{matrix} 1 \\ y_{10} \\ y_{01} \\ y_{20} \\ y_{11} \\ y_{02} \end{matrix} & \begin{matrix} y_{10} \\ y_{20} \\ y_{11} \\ y_{30} \\ y_{21} \\ y_{12} \end{matrix} & \begin{matrix} y_{01} \\ y_{11} \\ y_{02} \\ y_{21} \\ y_{12} \\ y_{03} \end{matrix} & \begin{matrix} y_{20} \\ y_{30} \\ y_{21} \\ y_{40} \\ y_{31} \\ y_{22} \end{matrix} & \begin{matrix} y_{11} \\ y_{21} \\ y_{12} \\ y_{31} \\ y_{22} \\ y_{13} \end{matrix} & \begin{matrix} y_{02} \\ y_{12} \\ y_{03} \\ y_{22} \\ y_{13} \\ y_{04} \end{matrix} \end{matrix}.$$

The matrix $\mathbf{M}_d(\mu)$ is positive semidefinite for all $d \in \mathbb{N}$. Indeed, for any $\mathbf{p} \in \mathbb{R}^{s(d)}$, let $P \in \mathbb{R}[X]_d$ be the polynomial with vector of coefficients \mathbf{p} , then $\mathbf{p}^T \mathbf{M}_d(\mu) \mathbf{p} = \int_{\mathbb{R}^p} P(\mathbf{x})^2 d\mu(\mathbf{x}) \geq 0$. We also have the identity $\mathbf{M}_d(\mu) = \int_{\mathbb{R}^p} \mathbf{v}_d(\mathbf{x}) \mathbf{v}_d(\mathbf{x})^T d\mu(\mathbf{x})$ where the integral is understood elementwise.

Sum of squares (SOS)

We denote by $\Sigma[X] \subset \mathbb{R}[X]$ (resp. $\Sigma[X]_d \subset \mathbb{R}[X]_d$), the set of polynomials (resp. polynomials of degree at most d) which can be written as a sum of squares of polynomials. Let $P \in \mathbb{R}[X]_{2m}$ for some $m \in \mathbb{N}$, then P belongs to $\Sigma[X]_{2m}$ if there exists a finite $J \subset \mathbb{N}$ and a family of polynomials $P_j \in \mathbb{R}[X]_m, j \in J$, such that $P = \sum_{j \in J} P_j^2$. It is obvious that in sum of squares polynomials are always nonnegative. A further interesting property is that this class of polynomials is connected with positive semidefiniteness. Indeed, P belongs to $\Sigma[X]_{2m}$ if and only if:

$$\exists Q \in \mathbb{R}^{s(m) \times s(m)}, Q = Q^T, Q \geq 0, P(\mathbf{x}) = \mathbf{v}_m(\mathbf{x})^T Q \mathbf{v}_m(\mathbf{x}), \forall \mathbf{x} \in \mathbb{R}^p. \quad (2.1)$$

As a consequence, every real symmetric positive semidefinite matrix $Q \in \mathbb{R}^{s(m) \times s(m)}$ defines a polynomial in $\Sigma[X]_{2m}$ by using the representation (2.1).

Orthonormal polynomials

We define a classical [16, 40] family of orthonormal polynomials, $\{P_\alpha\}_{\alpha \in \mathbb{N}_d^p}$ ordered according to \leq_{gl} , which satisfies for all $\alpha \in \mathbb{N}_d^p$:

$$\langle P_\alpha, P_\beta \rangle_\mu = \delta_{\alpha=\beta}, \quad \langle P_\alpha, X^\beta \rangle_\mu = 0, \text{ if } \beta <_{gl} \alpha, \langle P_\alpha, X^\alpha \rangle_\mu > 0. \quad (2.2)$$

Existence and uniqueness of such a family are guaranteed by the Gram-Schmidt orthonormalization process following the \leq_{gl} ordering on monomials and by the positivity of the moment matrix (see for instance [16] Theorem 3.1.11).

Let $\mathbf{D}_d(\mu)$ be the lower triangular matrix of which rows are the coefficients of the polynomials P_α defined in (2.2) ordered by \leq_{gl} . It can be shown that $\mathbf{D}_d(\mu) = \mathbf{L}_d(\mu)^{-T}$, where $\mathbf{L}_d(\mu)$ is the Cholesky factorization of $\mathbf{M}_d(\mu)$. Furthermore, there is a direct relation with the inverse moment matrix as $\mathbf{M}_d(\mu)^{-1} = \mathbf{D}_d(\mu)^T \mathbf{D}_d(\mu)$ ([20] Proof of Theorem 3.1).

3 The Christoffel function and its empirical counterpart

3.1 The Christoffel function

Let μ be a finite Borel measure on \mathbb{R}^p with all moments finite and such that its moment matrix $\mathbf{M}_d(\mu)$ is positive definite for every $d = 0, 1, \dots$. For every d , define the function $\kappa_{\mu,d} : \mathbb{R}^p \times \mathbb{R}^p \rightarrow \mathbb{R}$ by:

$$(\mathbf{x}, \mathbf{y}) \mapsto \kappa_{\mu,d}(\mathbf{x}, \mathbf{y}) := \sum_{\alpha \in \mathbb{N}_d^p} P_\alpha(\mathbf{x}) P_\alpha(\mathbf{y}) = \mathbf{v}_d(\mathbf{x})^T \mathbf{M}_d(\mu)^{-1} \mathbf{v}_d(\mathbf{y}), \quad (3.1)$$

where the family of polynomials $\{P_\alpha\}_{\alpha \in \mathbb{N}_d^p}$ is defined in (2.2) and the last equality follows from properties of this family (see also [23]). The kernel $(\mathbf{x}, \mathbf{y}) \mapsto K(\mathbf{x}, \mathbf{y}) := \sum_{\alpha \in \mathbb{N}^p} P_\alpha(\mathbf{x}) P_\alpha(\mathbf{y})$ is a reproducing kernel on $L_2(\mu)$ because:

$$P_\alpha(X) = \int K(X, \mathbf{y}) P_\alpha(\mathbf{y}) d\mu(\mathbf{y}), \quad \forall \alpha \in \mathbb{N}^p,$$

that is, the (P_α) are eigenvectors of the associated operator on $L_2(\mu)$, and so:

$$p(X) = \int K(X, \mathbf{y}) p(\mathbf{y}) d\mu(\mathbf{y}), \quad \forall p \in \mathbb{R}[X].$$

The function $\mathbf{x} \mapsto \Lambda_{\mu,d}(\mathbf{x}) := \kappa_{\mu,d}(\mathbf{x}, \mathbf{x})^{-1}$ is called the *Christoffel function* associated with μ and $d \in \mathbb{N}$. The following result states a fundamental extremal property of the Christoffel function.

Theorem 3.1 (see e.g. [16, 30]) *Let $\xi \in \mathbb{R}^p$ be fixed, arbitrary. Then,*

$$\Lambda_{\mu,d}(\xi) = \min_{P \in \mathbb{R}[X]_d} \left\{ \int_{\mathbb{R}^p} P(\mathbf{x})^2 d\mu(\mathbf{x}) : P(\xi) = 1 \right\}. \quad (3.2)$$

The Christoffel function plays an important role in orthogonal polynomials and the theory of interpolation and approximation (see e.g. [16, 40]). One is particularly interested in the asymptotics of the normalized Christoffel function $\mathbf{x} \mapsto s(d)\Lambda_{\mu,d}(\mathbf{x})$ as $d \rightarrow \infty$. The subject has a very long history in the univariate case (see [30] for a detailed historical account prior to the 1980s). The first quantitative asymptotic result was given in [27] and was latter improved by [28] and [41]. In the multivariate setting, precise results are known in some particular cases such as balls, spheres, and simplices [6, 7, 22, 45, 46] but much remains to be done for the general multivariate case. A typical example of asymptotic result is given under quite general (and technical) conditions in [21, 22]. This work shows that, as $d \rightarrow \infty$, the limit of the ratio, $\frac{\Lambda_{\nu,d}}{\Lambda_{\mu,d}}$, of two Christoffel functions associated to two mutually absolutely continuous measures μ and ν , converges to the density $\frac{d\nu}{d\mu}(\mathbf{x})$ on the interior of their common support.

Remark 3.2 Notice that Theorem 3.1 also provides a method to compute the numerical value $\Lambda_{\mu,d}(\mathbf{x})$ for $\mathbf{x} \in \mathbb{R}^n$, fixed, arbitrary. Then, indeed, (3.2) is a convex

quadratic programming problem which can be solved efficiently, even in high dimension using first-order methods such as projected gradient descent and its stochastic variants. This is particularly interesting when the nonsingular moment matrix $\mathbf{M}_d(\mu)$ is large.

We next provide additional insights on Theorem 3.1 and solutions of (3.2).

Theorem 3.3 *For any $\xi \in \mathbb{R}^p$, the optimization problem in (3.2) is convex with a unique optimal solution $P_d^* \in \mathbb{R}[X]_d$ defined by:*

$$P_d^*(X) = \frac{\kappa_{\mu,d}(X, \xi)}{\kappa_{\mu,d}(\xi, \xi)} = \Lambda_{\mu,d}(\xi) \kappa_{\mu,d}(X, \xi). \tag{3.3}$$

In addition,

$$\Lambda_{\mu,d}(\xi) = \int P_d^*(\mathbf{x})^2 d\mu(\mathbf{x}) = \int P_d^*(\mathbf{x}) d\mu(\mathbf{x}) \tag{3.4}$$

$$\Lambda_{\mu,d}(\xi) \xi^\alpha = \int \mathbf{x}^\alpha P_d^*(\mathbf{x}) d\mu(\mathbf{x}) \quad \alpha \in \mathbb{N}_d^p. \tag{3.5}$$

The proof is postponed to Section 6. Interestingly, each of the orthonormal polynomials $(P_\alpha)_{\alpha \in \mathbb{N}^p}$ also satisfies an important and well-known extremality property.

Theorem 3.4 (see, e.g., [16]) *Let $\alpha \in \mathbb{N}^p$ be fixed, arbitrary and let $d = |\alpha|$. Then up to a multiplicative positive constant, P_α is the unique optimal solution of:*

$$\min_{P \in \mathbb{R}[\mathbf{x}]_d} \left\{ \int P^2(\mathbf{x}) d\mu(\mathbf{x}) : P(\mathbf{x}) = \mathbf{x}^\alpha + \sum_{\beta <_{gl} \alpha} \theta_\beta \mathbf{x}^\beta \text{ for some } \{\theta_\beta\}_{\beta <_{gl} \alpha} \right\}. \tag{3.6}$$

Finally, we highlight the following important property which will be useful in the sequel.

Theorem 3.5 (See, e.g., [23]) *$\Lambda_{\mu,d}$ is invariant by change of polynomial basis \mathbf{v}_d , change of the origin of \mathbb{R}^p , or change of basis in \mathbb{R}^p .*

Remark 3.6 All these statements can be deduced from identity (3.1). Indeed, we have, for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^p$:

$$\begin{aligned} \mathbf{v}_d(\mathbf{x})^T \mathbf{M}_d(\mu)^{-1} \mathbf{v}_d(\mathbf{y}) &= (A \mathbf{v}_d(\mathbf{x}))^T \left(A \mathbf{M}_d(\mu) A^T \right)^{-1} (A \mathbf{v}_d(\mathbf{y})) \\ &= (A \mathbf{v}_d(\mathbf{x}))^T \left(\int_{\mathbb{R}^p} (A \mathbf{v}_d(\mathbf{z})) (A \mathbf{v}_d(\mathbf{z}))^T d\mu(\mathbf{z}) \right)^{-1} (A \mathbf{v}_d(\mathbf{y})) \end{aligned}$$

for any invertible matrix A of suitable size. All the proposed transformations induce a change of basis of polynomials up to degree d which can be represented by such an A .

3.2 When μ is the Lebesgue measure

In this section, we consider the important case of the Lebesgue measure on a compact set $S \subset \mathbb{R}^p$ such that $\text{cl}(\text{int}(S)) = S$. It is known that in this case the Christoffel function encodes information on the set S (see for example the discussion in Section 3.1). In particular, the scaled Christoffel function remains positive on the interior of S . We push this idea further and present a new result asserting that it is possible to recover the set S with strong asymptotic guaranties by carefully thresholding the corresponding scaled Christoffel function.

For any measurable set A , denote by μ_A the uniform probability measure on A , that is $\mu_A = \lambda_A/\lambda(A)$ where λ is the Lebesgue measure and λ_A the measure consisting of the restriction of Lebesgue measure to A which is defined by $\lambda_A(A') = \lambda_{A \cap A'}$ for any measurable set A' .

Threshold and asymptotics

The main idea is to use quantitative lower bounds on the scaled Christoffel function, $s(d)\Lambda_{\mu_S, d}$, on the interior of S (Lemma 6.2) and upper bounds outside S (Lemma 6.6). Recall that μ_S denotes the uniform measure on S . In combining these bounds, one proves the existence of a sequence of thresholds of the scaled Christoffel function which estimates S in a strongly consistent manner. Let us introduce the following notation and assumption.

- Assumption 3.7** (a) $S \subset \mathbb{R}^p$ is a compact set such that $\text{cl}(\text{int}(S)) = S$.
 (b) The sequence $(\delta_k)_{k \in \mathbb{N}}$ is a decreasing sequence of positive numbers converging to 0. For every $k \in \mathbb{N}$, let d_k be the smallest integer such that:

$$2^{3 - \frac{\delta_k d_k}{\delta_k + \text{diam}(S)}} d_k^p \left(\frac{e}{p}\right)^p \exp\left(\frac{p^2}{d_k}\right) \leq \alpha_k \tag{3.7}$$

where $\text{diam}(S)$ denotes the diameter of the set S , and:

$$\alpha_k := \frac{\delta_k^p \omega_p}{\lambda(S)} \frac{(d_k + 1)(d_k + 2)(d_k + 3)}{(d_k + p + 1)(d_k + p + 2)(2d_k + p + 6)}$$

Remark 3.8 (On Assumption 3.7)

- d_k is well defined. Indeed, since δ_k is positive, the left-hand side of (3.7) goes to 0 as $k \rightarrow \infty$ while the right-hand side remains bounded for increasing values of d_k .
- From the definition of d_k and the fact that δ_k is decreasing, the sequence $\{d_k\}_{k \in \mathbb{N}}$ is nondecreasing. Indeed, in (3.7), the right-hand side is an increasing function of δ_k while the left-hand side is decreasing so that if (3.7) is satisfied for a certain value of d_k and δ_k , it is also satisfied with the same d_k and any value of $\delta \geq \delta_k$.
- Given $\{\delta_k\}_{k \in \mathbb{N}}$, computing d_k can be done recursively and only requires the knowledge of $\text{diam}(S)$ and $\lambda(S)$.

- A similar condition can be enforced if only upper bounds on $\text{diam}(S)$ and on $\lambda(S)$ are available. In this case, replace these quantities by their upper bounds in (3.7) to obtain a similar result.

We are now ready to state the first main result of this section whose proof is postponed to Section 6 for the sake of clarity of exposition. Recall the definition of the Hausdorff distance $d_H(X, Y)$ between two subsets X, Y of \mathbb{R}^p :

$$d_H(X, Y) = \max \left\{ \sup_{\mathbf{x} \in X} \inf_{\mathbf{y} \in Y} \text{dist}(\mathbf{x}, \mathbf{y}), \sup_{\mathbf{y} \in Y} \inf_{\mathbf{x} \in X} \text{dist}(\mathbf{x}, \mathbf{y}) \right\}.$$

Theorem 3.9 *Let $S \subset \mathbb{R}^p$, $\{\delta_k\}_{k \in \mathbb{N}}$, $\{\alpha_k\}_{k \in \mathbb{N}}$ and $\{d_k\}_{k \in \mathbb{N}}$ satisfy Assumption 3.7. For every $k \in \mathbb{N}$, let $S_k \subset \mathbb{R}^p$ be the set defined by:*

$$S_k := \{ \mathbf{x} \in \mathbb{R}^p : s(d_k) \Lambda_{\mu_S, d_k}(\mathbf{x}) \geq \alpha_k \}.$$

Then, as $k \rightarrow \infty$:

$$\begin{aligned} d_H(S_k, S) &\rightarrow 0 \\ d_H(\partial S_k, \partial S) &\rightarrow 0. \end{aligned}$$

Remark 3.10 The relevance of Hausdorff distance and the notion of distance between topological boundaries are discussed in [12] and [39].

3.3 Extension to more general probability measures

Theorem 3.9 can easily be extended to probability measures that are more general than uniform distributions, in which case we consider the following alternative assumption.

Assumption 3.11 (a) $S \subset \mathbb{R}^p$ is a compact set such that $\text{cl}(\text{int}(S)) = S$.

(b) The function $w : \text{int}(S) \rightarrow [w_-, +\infty)$ is integrable on $\text{int}(S)$ with $w_- > 0$. The measure μ is such that for any measurable set A , $\mu(A) = \int_{A \cap S} w(\mathbf{x}) d\mathbf{x}$ and $\mu(S) = 1$. $\{\delta_k\}_{k \in \mathbb{N}}$ is a decreasing sequence of positive numbers which converges to 0. For every $k \in \mathbb{N}$, let d_k be the smallest integer such that:

$$2^{3 - \frac{\delta_k d_k}{\delta_k + \text{diam}(S)}} d_k^p \left(\frac{e}{p}\right)^p \exp\left(\frac{p^2}{d_k}\right) \leq \alpha_k \tag{3.8}$$

where:

$$\alpha_k := w_- \delta_k^p \omega_p \frac{(d_k + 1)(d_k + 2)(d_k + 3)}{(d_k + p + 1)(d_k + p + 2)(2d_k + p + 6)}.$$

Under Assumption 3.11, we obtain the following analogue of Theorem 3.9.

Theorem 3.12 *Let $S \subset \mathbb{R}^p$, $w : S \rightarrow [w_-, +\infty)$, $\{\delta_k\}_{k \in \mathbb{N}}$, $\{\alpha_k\}_{k \in \mathbb{N}}$ and $\{d_k\}_{k \in \mathbb{N}}$ satisfy Assumption 3.11. For every $k \in \mathbb{N}$, let:*

$$S_k := \{ \mathbf{x} \in \mathbb{R}^p : s(d_k) \Lambda_{\mu_S, d_k}(\mathbf{x}) \geq \alpha_k \}.$$

Then, as $k \rightarrow \infty$:

$$\begin{aligned} d_H(S_k, S) &\rightarrow 0 \\ d_H(\partial S_k, \partial S) &\rightarrow 0. \end{aligned}$$

The proof is postponed to Section 6.

3.4 Discrete approximation via the empirical Christoffel function

In this section, μ is a probability measure on \mathbb{R}^p with compact support S . We focus on the statistical setting where information on μ is available only through a sample of points drawn independently from the given distribution μ . In this setting, for every $n \in \mathbb{N}$, let μ_n denote the empirical measure uniformly supported on an independent sample of n points distributed according to μ . It is worth emphasizing that in principle $\Lambda_{\mu_n, d}$ is easy to compute and requires the inversion of a square matrix of size $s(d)$ (see (3.1)). Note that the definition in (3.1) can only be used if the empirical moment matrix $\mathbf{M}_d(\mu_n)$ is invertible which is the case almost surely if $\mathbf{M}_d(\mu)$ is invertible and n is large enough. Alternatively, the numerical evaluation of $\Lambda_{\mu_n, d}(\mathbf{x})$ at $\mathbf{x} \in \mathbb{R}^n$, fixed arbitrary, reduces to solving the convex quadratic programming problem (3.2), which can in principle be done efficiently even in high dimension (see Remark 3.2).

Our second main result is for fixed $d \in \mathbb{N}$ and relates the population Christoffel function $\Lambda_{\mu, d}$ and its empirical version $\Lambda_{\mu_n, d}$, as n increases. We proceed by distinguishing what happens far from S and close to S . First, Lemma 6.8 ensures that both Christoffel functions associated with μ and μ_n vanish far from S so that the influence of this region can be neglected. Second, when closer to S , one remains in a compact set and the strong law of large numbers applies.

Theorem 3.13 *Let μ be a probability measure on \mathbb{R}^p with compact support. Let $\{X_i\}_{i \in \mathbb{N}}$ be a sequence of i.i.d. \mathbb{R}^p -valued random variables with common distribution μ . For $n = 1, 2, \dots$, define the (random) empirical probability measure $\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$. Then, for every $d \in \mathbb{N}$, $d > 0$, such that the moment matrix $\mathbf{M}_d(\mu)$ is invertible, it holds that:*

$$\sup_{\mathbf{x} \in \mathbb{R}^p} \{ |\Lambda_{\mu_n, d}(\mathbf{x}) - \Lambda_{\mu, d}(\mathbf{x})| \} \xrightarrow[n \rightarrow \infty]{a.s.} 0. \tag{3.9}$$

Equivalently:

$$\| \Lambda_{\mu_n, d} - \Lambda_{\mu, d} \|_\infty \xrightarrow[n \rightarrow \infty]{a.s.} 0. \tag{3.10}$$

where $\| \cdot \|_\infty$ denotes the usual “sup-norm.”

A detailed proof can be found in Section 6. Theorem 3.13 is a strong result which states a highly desirable property, namely that *almost surely* with respect to the random draw of the sample, the (random) function $\Lambda_{\mu_n, d}(\cdot)$ converges to $\Lambda_{\mu, d}(\cdot)$ *uniformly* in $\mathbf{x} \in \mathbb{R}^p$ as n increases. Since we manipulate polynomials, it can be checked that [12, Theorem 1] for general level sets can be applied in the setting of Christoffel level set estimation. We get the following consequence in terms of consistency of the boundary of plugin estimates for Christoffel level sets.

Theorem 3.14 *Let μ be a probability measure on \mathbb{R}^P with compact support. Let $\{X_i\}_{i \in \mathbb{N}}$ be a sequence of i.i.d. \mathbb{R}^P -valued random variables with common distribution μ . For $n = 1, 2, \dots$, define the (random) empirical probability measure $\mu_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$. Then, for every $d \in \mathbb{N}$, $d > 0$, such that the moment matrix $\mathbf{M}_d(\mu)$ is invertible and any $c \in (0, \sup_{\mathbf{x} \in \mathbb{R}^P} \{\Lambda_{\mu,d}(\mathbf{x})\})$, as n increases, it holds that:*

$$d_H(\partial L_n, \partial L) \xrightarrow{a.s.} 0, \quad (3.11)$$

where $L = \{\mathbf{x} \in \mathbb{R}^P, \Lambda_{\mu,d}(x) \geq c\}$ and $L_n = \{\mathbf{x} \in \mathbb{R}^P, \Lambda_{\mu_n,d}(x) \geq c\}$.

4 Applications

4.1 Rationale

In this section, we describe some applications for which properties of the Christoffel function prove to be very useful in a statistical context. We only consider the case of bounded support. A relevant property of the scaled Christoffel function is that it encodes information on the support and the density of a population measure μ :

- [28, 41] and [21] provide asymptotic results involving the density of the input measure.
- Theorem 3.9 provides asymptotic results related to the support of the input measure.

The support and density of a measure are of interest in many statistical applications. However, the aforementioned results are limited to population measures which are not accessible in a statistical setting. In the context of empirical Christoffel functions, Theorem 3.13 suggests that these properties still hold (at least in the limit of large number of samples) when one uses the empirical measure μ_n in place of the population measure μ . Combining these ideas suggests the use of the empirical Christoffel function in statistical applications such as (a) density estimation, (b) support inference, or (c) outlier detection. This is illustrated on simulated and real-world data and we compare the performance with well-established methods for the same purpose. Finally, we also describe another application, namely *inversion of affine shuffling*, whose links with statistics are less clear.

All results presented in this section are mainly for illustrative purposes. In particular, the choice of the degree d as a function of the sample size n was done empirically and a precise quantitative analysis is a topic of future research beyond the scope of the present paper.

4.2 Density estimation

Most asymptotic results regarding the scaled Christoffel function suggest that the limiting behavior involves the product of a boundary effect term and a density term. Hence, if one knows both the Christoffel function and the boundary effect term, one has access to the density term. Unfortunately, this boundary term is only known in

specific situations, the most typical example being the Euclidean ball. Hence, in the present state of knowledge, one of the following is assumed to hold true.

- The support of the population measure μ is S and $\lim_{d \rightarrow \infty} s(d)\Lambda_{\lambda_S, d}$ exists (possibly unknown).
- The support is unknown but contains a set S with the same property as above. In this case, we consider the restriction of the population measure μ to S . Note that a sample from the restriction is easily obtained from a sample from μ by rejection.

In both cases, assuming that μ has a density h on S , it is expected that the ratio $\frac{\Lambda_{\mu, d}}{\Lambda_{\lambda_S, d}}$ or $\frac{s(d)\Lambda_{\mu, d}}{\lim_{d \rightarrow \infty} s(d)\Lambda_{\lambda_S, d}}$ converges to h . An example of such a result in the univariate setting is the following.

Theorem 4.1 (Theorem 5 [28]) *Suppose that μ is supported on $[-1, 1]$ with density $h \geq a > 0$. Then, for almost every $x \in [-1, 1]$:*

$$\lim_{d \rightarrow \infty} d\Lambda_{\mu, d}(x) \rightarrow \pi h(x)\sqrt{1 - x^2}.$$

Extensions include [41] for general support and [21] for the multivariate setting. Combining Theorems 4.1 and 3.13 suggests that the empirical Christoffel function can be used for density estimation. For illustration purposes, we set μ to be the restriction of a Gaussian to $[-1, 1]$. We perform the following experiment for given $n, d \in \mathbb{N}$.

- Generate $x_1, \dots, x_n \in [-1, 1]$ sampled independently from μ .
- Compute and plot $x \rightarrow \frac{\Lambda_{\mu_n, d}(x)}{\Lambda_{\lambda_S, d}(x)}$. Note that $\Lambda_{\lambda_S, d}$ is easily derived from the moments of the uniform distribution on $[-1, 1]$.

The result is presented in Fig. 1 and a comparison is given with a classical technique, kernel density estimation [31, 37] with the Gaussian kernel. The result suggests that empirical Christoffel-based density estimation is competitive with kernel density estimation in this setting. It is worth noticing how simple the methodology is with a single parameter to tune.

4.3 Support inference

Combining Theorems 3.9 and 3.13 suggests that one may recover the unknown support of a population measure μ from n independent samples by thresholding the scaled empirical Christoffel function. In this section, we set μ to be the uniform probability measure over a star-shaped domain in \mathbb{R}^2 (see Fig. 2). For different values of the degree d and sample size n , we plot in Fig. 2 the corresponding sample and the associated level set $\{\mathbf{x} \in \mathbb{R}^2 : s(d)\Lambda_{\mu_n, d}(\mathbf{x}) = \alpha(\delta)\}$, where $\delta = 0.5$ and α is given in Assumption 3.7.

The results displayed in Fig. 2 show that for well-chosen values of d and with α as in Assumption 3.7, the support of the population measure is rather well approximated from a finite independent sample. The results even suggest that a careful tuning of the degree d and the threshold level set α allows approximating the support extremely

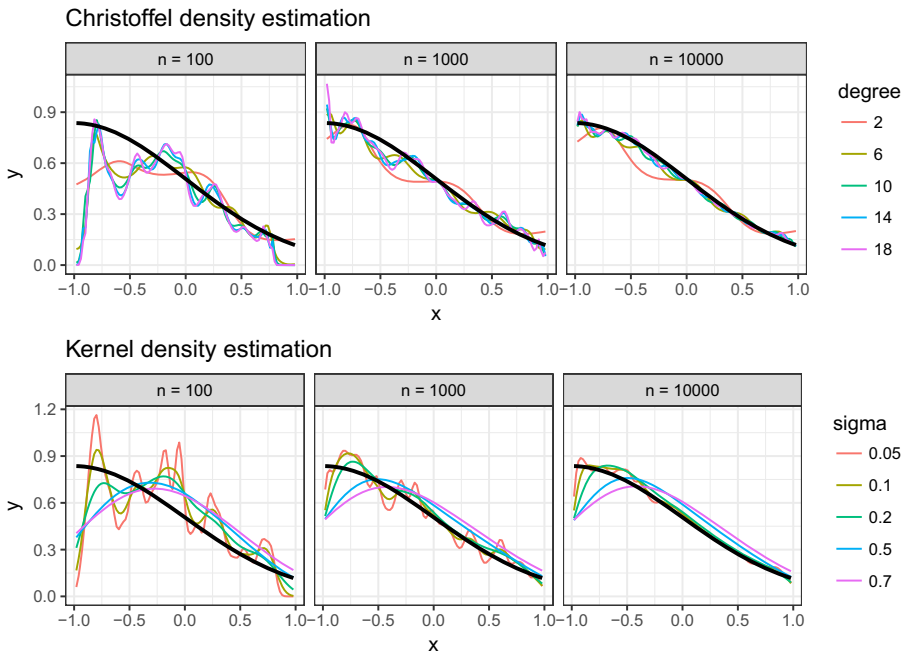


Fig. 1 Comparison of Christoffel and kernel density estimation with Gaussian kernel. The same samples are used in both cases. We vary the sample size n , the degree d for the Christoffel function, and the scale parameter σ for the Gaussian kernel. The black curve shows the population density

well for larger sample sizes. Of course, the degree d should be chosen to avoid a form of overfitting as the results suggest for small sample sizes and large values of d . A precise analysis of this phenomenon is a topic of future research.

4.4 Outlier detection

In [23], we suggested that the empirical Christoffel function could be used for the purpose of detecting outliers and the claim was supported by some numerical experiments. The rationale for this is that the empirical Christoffel function encodes information about the population density and outliers can be seen as samples from low-density areas. We follow the same line and consider the network intrusion detection task described in [43] based on the KDD cup 99 dataset [24]. Following the pre-processing described in [23, 43], we build up five datasets consisting of network connections represented by labeled vectors in \mathbb{R}^3 where each label indicates whether the connection was an attack or not.

Dataset	http	smtp	ftp-data	ftp	Other
Number of examples	567498	95156	30464	4091	5858
Proportions of attacks	0.004	0.0003	0.023	0.077	0.016

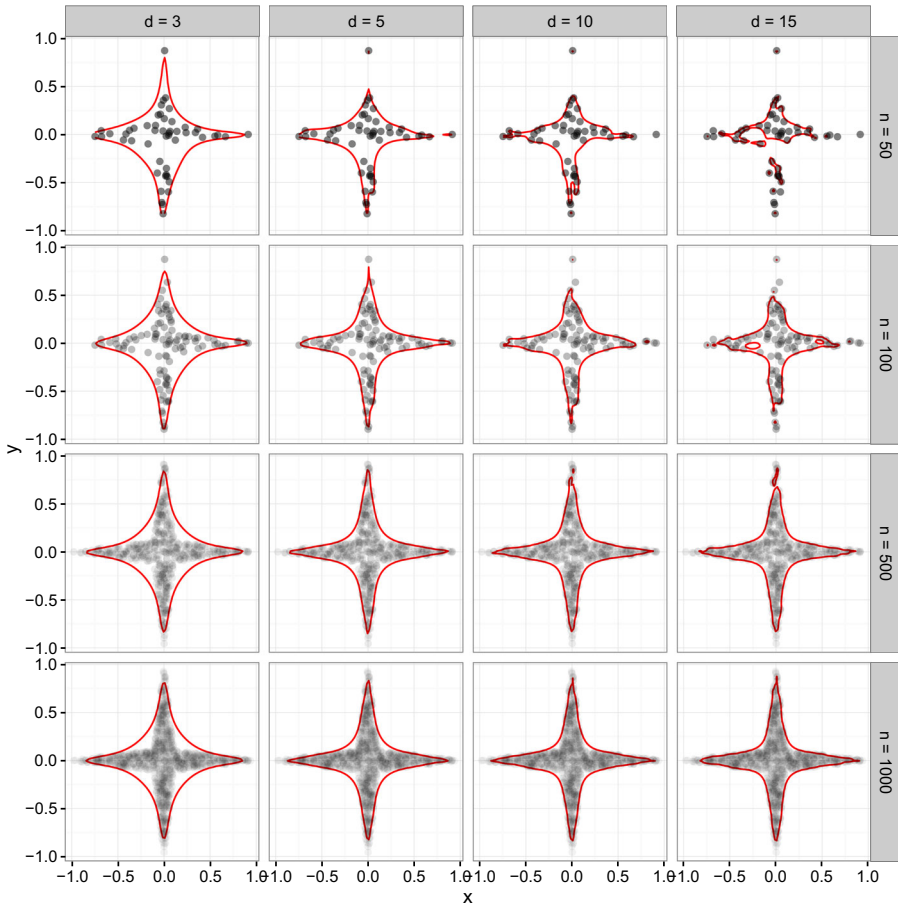


Fig. 2 Finite sample from the uniform measure over a star-shaped domain in \mathbb{R}^2 . For each value of n and d , the red line represents the sublevel set $\{x \in \mathbb{R}^2, s(d)\Lambda_{\mu_n, d}(x) = \alpha(\delta)\}$, where $\delta = 0.5$ and α is given in Assumption 3.7

All the experiments described in this section are performed on the “Other” dataset which is the most heterogeneous. The main task is to recover attacks from the collection of points in \mathbb{R}^3 , ignoring the labels, and then compare the predictions with the ground truth (given by the labels). We compare different methods, each of them assigning a score to an individual: the higher the score, the more likely the individual is to be an outlier, or an attack. The metric that we use to compare different methods is the area under the Precision Recall curve (AUPR) (see for example [13]). We compare three different methods, each of them producing a score reflecting some degree of outlyingness.

- Empirical Christoffel function.
- Kernel density estimation [31, 37] with Gaussian kernel. The value of the density estimated at each datapoint is used as an outlyingness score.

- One-class SVM [38] with Gaussian kernel. The value of the estimated decision function at each datapoint is used as an outlyingness score. We used the implementation provided in the kernlab package [47].

The first two methods involve only a single parameter while the last method requires two parameters to be tuned. The results are given in Table 1 and the corresponding curves can be found in Section I. Table 1 suggests that one-class SVM and the empirical Christoffel perform similarly and clearly outperform the kernel density estimation approach. It is worth noticing here that the one-class SVM provides slightly better performances but requires a precise tuning of the second parameter.

4.5 Inversion of affine shuffling

This last application has fewer connections with statistics. Suppose that we are given two matrices $\mathbf{X} \in \mathbb{R}^{p \times n}$ and $\mathbf{X}' \in \mathbb{R}^{p \times n}$. Furthermore, we know that there exists an invertible affine mapping $\mathcal{A}: \mathbb{R}^p \rightarrow \mathbb{R}^p$ such that after a potential permutation of the columns, \mathcal{A} defines a bijection between the columns of \mathbf{X}' and those of \mathbf{X} . The problem is to recover the correspondence between the columns of \mathbf{X} and the columns of \mathbf{X}' , hence the name “affine shuffling inversion.” Note that the columns may be shuffled in an arbitrary way and therefore the matching problem is not trivial. In this setting, we can use the affine invariance property of the Christoffel function described in Theorem 3.5. This is based on the two following observations.

- The Christoffel function only depends on the empirical moments and hence is not sensitive to reshuffling of the columns.

Table 1 AUPR scores ($\times 100$) for the network intrusion detection task for the three different methods considered in this paper

1SVM							Christoffel		KDE	
$\sigma \backslash \nu$	0.005	0.01	0.02	0.05	0.1	0.2	d	AUPR	σ	AUPR
0.01	10	17	17	17	15	11	1	8	0.01	8
0.02	2	17	18	17	15	12	2	18	0.02	1
0.05	8	1	14	18	15	11	3	18	0.05	13
0.1	9	8	12	17	14	11	4	16	0.1	13
0.2	7	9	8	17	14	13	5	15	0.2	12
0.5	3	5	9	15	17	16	6	13	0.5	5
1	3	6	9	14	19	18			1	4
2	4	4	5	1	18	18				
5	4	3	4	9	12	16				
10	5	4	4	7	10	15				

1SVM corresponds to one-class SVM with Gaussian kernel and varying kernel scale parameter σ and SVM parameter ν . Christoffel corresponds to the empirical Christoffel function with varying degree d . KDE corresponds to kernel density estimation with Gaussian kernel and varying scale parameter σ . The best scores are highlighted in boldface font




Fig. 3 The affine matching procedure based on the empirical Christoffel function. On the left, two datasets of points in \mathbb{R}^2 : the second one is the image of the first one by an affine transformation. The colors indicate the correspondence between the two clouds of points which have been recovered by matching the corresponding empirical Christoffel functions. The matching of these values is illustrated on the right with a quantile–quantile plot of the empirical Christoffel function values for each dataset. We applied a log transformation for readability and the first diagonal is represented

- Working with the affine image amounts to performing a change of basis and a change of origin. By Theorem 3.5, the evaluation of the Christoffel function does not change.

This suggests the following procedure.

- Compute $\Lambda_{\mathbf{X},d}$ and $\Lambda_{\mathbf{X}',d}$ the Christoffel functions associated to the columns of \mathbf{X} and the columns of \mathbf{X}' respectively.
- Set $A \in \mathbb{R}^n$ to be the vector with $\Lambda_{\mathbf{X},d}(\mathbf{X}_i)$ as the i th entry where \mathbf{X}_i is the i th column of \mathbf{X} . Set A' similarly.
- Match the values in A to the values in A' according to their rank.

The proposed procedure defines a unique permutation between columns of \mathbf{X} and columns of \mathbf{X}' when there are no ties in the vectors A and A' . In this case, Theorem 3.5 ensures that we have found the correct correspondence. In case of ties, the procedure does not allow eliciting completely the correspondence matching. Overall, the method is not guaranteed to work but allows treating simple cases easily. Investigating the robustness of this procedure to noise or to matching misspecification is the subject of future research.

An illustration is given in Fig. 3 where a moon-shaped cloud of points in \mathbb{R}^2 is deformed by an affine transformation and the matching between the points between the two clouds is recovered by matching the corresponding Christoffel function values. The correspondence between Christoffel function values is illustrated on a quantile–quantile plot.

5 Conclusion

In this paper, we have investigated the potential of the empirical Christoffel function for some applications in statistics and machine learning. This question led us to

investigate its theoretical properties as well as potential paths toward applications, mostly in a statistical framework.

On the theoretical side, we proposed two main contributions. The first one provides an explicit thresholding scheme which allows the use of the Christoffel function to recover the support of a measure with strong asymptotic guarantees. Although this property finds its root in the long history of results regarding asymptotic properties of the Christoffel function, we have provided a systematic way to tune the threshold and the degree to ensure strong convergence guarantees. The second main contribution relates the empirical Christoffel function to its population counterpart in the limit of large samples. This type of results is new and paves the way toward a much more precise understanding of relations between these two objects in a small sample setting.

On the practical side, we have illustrated the relevance of the Christoffel function as a practical tool in a machine learning context. In particular, simulations and experiments on real-world data support our claim that the empirical Christoffel function is potentially useful for density estimation and support inference and outlier detection. Finally, in another application outside the statistical framework (detection of affine matching between two clouds of points), we have again illustrated the potential of the Christoffel function as a tool in shape recognition and shape comparison.

Both theory and applications suggest a broad research program. As already mentioned, an important issue is to quantify the deviation of the empirical Christoffel function from its population counterpart in a finite sample setting. Results in this direction could have both theoretical and practical impacts and would compare more accurately the performance of Christoffel-based approaches with state-of-the-art methods. Furthermore, the use of the Christoffel function in a statistical framework raises questions specific to each application considered in this paper and will be the subject of future investigations. Finally, there are still important open questions on the Christoffel function itself and works in the line of [5] are of great interest to address applications in statistics.

Acknowledgements The research of the first author was funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement 666981 TAMING).

Appendix

Precision recall curves from Section 4.4

This section displays the curves from which the AUPR scores were measured in Section 4.4. Christoffel function and kernel density estimation are presented in Fig. 4 and the one-class SVM is presented in Fig. 5. A detailed discussion on the experiment is given in Section 4.4.




Fig. 4 Precision recall curves for the network intrusion detection task. Left: Christoffel function with varying degree d . Right: kernel density estimation with Gaussian kernel and varying scale parameter σ

Proof of Theorem 3.3

Proof In the optimization problem (3.2), the objective function $P \mapsto \int P^2 d\mu$ is strongly convex in the vector of coefficients of P because:

$$\int P^2 d\mu = P^T \mathbf{M}_d(\mu) P \quad \text{and} \quad \mathbf{M}_d(\mu) \succ 0,$$

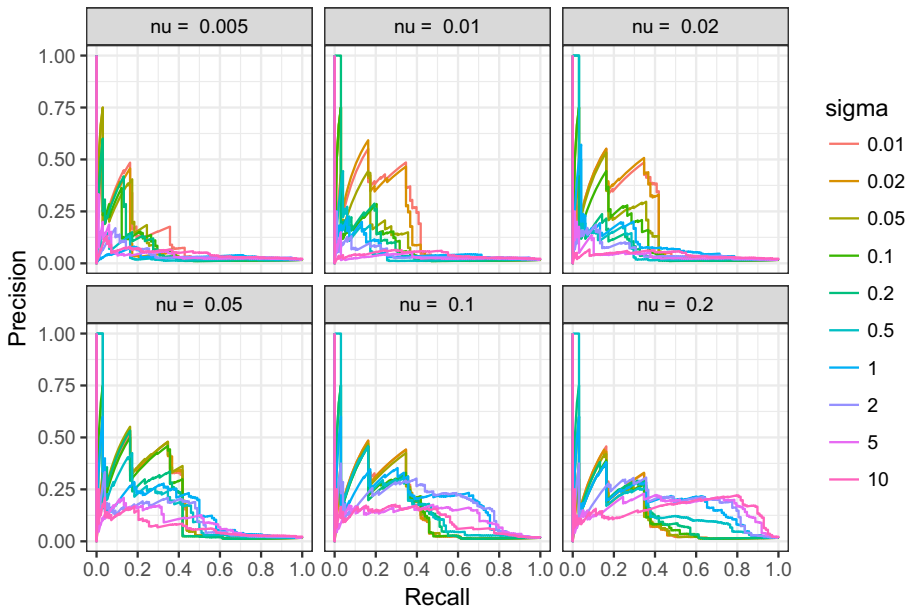


Fig. 5 Precision recall curves for the network intrusion detection task. The method used is the one-class SVM with a Gaussian kernel. We vary the scale parameter σ and the SVM parameter ν . We used the SVM solver of the package [47]

and therefore (3.2) reads $\min\{P^T \mathbf{M}_d(\mu) P : P^T \mathbf{v}_d(\xi) = 1\}$, which is a convex optimization problem with a strongly convex objective function. Slater’s condition holds (only one linear equality constraint) and so the Karush-Kuhn-Tucker (KKT) optimality conditions are both necessary and sufficient. At an optimality solution P_d^* , they read:

$$P_d^*(\xi) = 1; \quad 2\mathbf{M}_d(\mu) P_d^* = \theta \mathbf{v}_d(\xi),$$

for some scalar θ . Multiplying by $(P_d^*)^T$ yields:

$$2\kappa_{\mu,d}(\xi, \xi)^{-1} = 2(P_d^*)^T \mathbf{M}_d(\mu) P_d^* = \theta.$$

Hence, necessarily:

$$P_d^*(X) = \mathbf{v}_d(X)^T P_d^* = \frac{\theta}{2} \mathbf{v}_d(X)^T \mathbf{M}_d(\mu)^{-1} \mathbf{v}_d(\xi) = \frac{\kappa_{\mu,d}(X, \xi)}{\kappa_{\mu,d}(\xi, \xi)},$$

which is (3.3). Next, let $\mathbf{e}_\alpha \in \mathbb{R}^{s(d)}$ be the vector with null coordinates except the entry α which is 1. From the definition of the moment matrix $\mathbf{M}_d(\mu)$,

$$\mathbf{e}_\alpha^T \mathbf{M}_d(\mu) P_d^* = \int \mathbf{x}^\alpha P^*(\mathbf{z}) d\mu(\mathbf{z}) = \kappa_{\mu,d}(\xi, \xi)^{-1} \mathbf{e}_\alpha^T \mathbf{v}_d(\xi) = \kappa_{\mu,d}(\xi, \xi)^{-1} \xi^\alpha,$$

which is (3.5). In particular with $\alpha := 0$, we recover (3.4):

$$\int P_d^*(\mathbf{x}) d\mu(\mathbf{x}) = \kappa_{\mu,d}(\xi, \xi)^{-1} = \int P_d^*(\mathbf{x})^2 d\mu(\mathbf{x}).$$

□

Proof of Theorems 3.9 and 3.12

Lower bound on the Christoffel function inside S

We will heavily rely on results from [6] (note that similar results could be obtained on the box, see for example [44]). In particular, we have the following result.

Lemma 6.1 *We have for any $d \geq 2$*

$$\frac{\kappa_{\lambda_{\mathbf{B}},d}(0, 0)}{s(d)} \leq \frac{1}{\omega_p} \frac{(d + p + 1)(d + p + 2)}{(d + 1)(d + 2)} \left(1 + \frac{d + p + 3}{d + 3}\right)$$

Proof Combining Lemma 2 in [6] and the last equation of the proof of Lemma 3 in [6], we have:

$$\kappa_{\lambda_{\mathbf{B}},d}(0, 0) \leq \frac{1}{\omega_p} \left(\binom{p + d + 3}{p} + \binom{p + d + 2}{p} \right).$$

The result follows by using the expression given for $s(d)$ and simplifying factorial terms. □

From this result, we deduce the following bound.

Lemma 6.2 *Let $\delta > 0$ and $\mathbf{x} \in S$ such that $\text{dist}(\mathbf{x}, \partial S) \geq \delta$. Then,*

$$s(d)\Lambda_{\mu_S,d}(\mathbf{x}) \geq \frac{\delta^p \omega_p}{\lambda(S)} \frac{(d+1)(d+2)(d+3)}{(d+p+1)(d+p+2)(2d+p+6)}.$$

Proof Decompose the measure μ_S into the sum,

$$\mu_S = \frac{\lambda(S \setminus \mathbf{B}_\delta(\mathbf{x}))}{\lambda(S)} \mu_{S \setminus \mathbf{B}_\delta(\mathbf{x})} + \frac{\lambda(\mathbf{B}_\delta(\mathbf{x}))}{\lambda(S)} \mu_{\mathbf{B}_\delta(\mathbf{x})}.$$

Hence, by monotonicity of the Christoffel function with respect to addition and closure under multiplication by a positive term (this follows directly from Theorem 3.1), we have:

$$\Lambda_{\mu_S,d}(\mathbf{x}) \geq \frac{\lambda(\mathbf{B}_\delta(\mathbf{x}))}{\lambda(S)} \Lambda_{\mu_{\mathbf{B}_\delta(\mathbf{x})},d}(\mathbf{x}). \tag{6.1}$$

Next, by affine invariance of the Christoffel function (Theorem 3.5):

$$\Lambda_{\mu_{\mathbf{B}_\delta(\mathbf{x})},d}(\mathbf{x}) = \Lambda_{\mu_{\mathbf{B}},d}(0) = \frac{1}{\lambda(\mathbf{B})} \Lambda_{\lambda_{\mathbf{B}},d}(0) = \frac{1}{\lambda(\mathbf{B})} \frac{1}{\kappa_{\lambda_{\mathbf{B}},d}(0,0)}, \tag{6.2}$$

where \mathbf{B} is the unit Euclidean ball in \mathbb{R}^p . The result follows by combining (6.1), (6.2), Lemma 6.1 and the fact that $\frac{\lambda(\mathbf{B}_\delta(\mathbf{x}))}{\lambda(\mathbf{B})} = \delta^p$. \square

Upper bound on the Christoffel function outside S

We next exhibit an upper bound on the Christoffel function outside of S . We first provide a useful quantitative refinement of the ‘‘Needle polynomial’’ introduced in [21].

Lemma 6.3 *For any $d \in \mathbb{N}$, $d > 0$, and any $\delta \in (0, 1)$, there exists a p -variate polynomial of degree $2d$, q , such that:*

$$q(\mathbf{0}) = 1; \quad -1 \leq q \leq 1, \text{ on } \mathbf{B}; \quad |q| \leq 2^{1-\delta d} \text{ on } \mathbf{B} \setminus \mathbf{B}_\delta(\mathbf{x}).$$

Proof Let r be the univariate polynomial of degree $2d$, defined by:

$$r: t \rightarrow \frac{T_d(1 + \delta^2 - t^2)}{T_d(1 + \delta^2)},$$

where T_d is the Chebyshev polynomial of the first kind. We have

$$r(0) = 1. \tag{6.3}$$

Furthermore, for $t \in [-1, 1]$, we have $0 \leq 1 + \delta^2 - t^2 \leq 1 + \delta^2$. T_d has absolute value less than 1 on $[-1, 1]$ and is increasing on $[1, \infty)$ with $T_d(1) = 1$, so for $t \in [-1, 1]$,

$$-1 \leq r(t) \leq 1. \tag{6.4}$$

For $|t| \in [\delta, 1]$, we have $\delta^2 \leq 1 + \delta^2 - t^2 \leq 1$, so

$$|r(t)| \leq \frac{1}{T_d(1 + \delta^2)}. \tag{6.5}$$

Let us bound the last quantity. Recall that for $t \geq 1$, we have the following explicit expression:

$$T_d(t) = \frac{1}{2} \left((t + \sqrt{t^2 - 1})^d + (t + \sqrt{t^2 - 1})^{-d} \right).$$

We have $1 + \delta^2 + \sqrt{(1 + \delta^2)^2 - 1} \geq 1 + \sqrt{2}\delta$, which leads to

$$\begin{aligned} T_d(1 + \delta^2) &\geq \frac{1}{2} (1 + \sqrt{2}\delta)^d \\ &= \frac{1}{2} \exp(\log(1 + \sqrt{2}\delta) d) \\ &\geq \frac{1}{2} \exp(\log(1 + \sqrt{2})\delta d) \\ &\geq 2^{\delta d - 1}, \end{aligned} \tag{6.6}$$

where we have used concavity of the log and the fact that $1 + \sqrt{2} \geq 2$. It follows by combining (6.3), (6.4), (6.5), and (6.6), that $q: y \rightarrow r(\|y - x\|_2)$ satisfies the claimed properties. \square

We recall the following well-known bound for the factorial taken from [36].

Lemma 6.4 ([36]) *For any $n \in \mathbb{N}$, we have:*

$$\exp\left(\frac{1}{12n + 1}\right) \leq \frac{n!}{\sqrt{2\pi n} n^n \exp(-n)} \leq \exp\left(\frac{1}{12n}\right).$$

We deduce the following Lemma.

Lemma 6.5 *For any $d \in \mathbb{N}$, $d > 0$, we have:*

$$\binom{p + d}{d} \leq d^p \left(\frac{e}{p}\right)^p \exp\left(\frac{p^2}{d}\right)$$

Proof This follows from a direct computation using Lemma 6.4.

$$\begin{aligned} \binom{p + d}{d} &= \frac{(p + d)!}{p!d!} \\ &\leq \frac{\exp\left(\frac{1}{24}\right)}{\sqrt{2\pi}} \sqrt{\frac{p + d}{pd}} \frac{(p + d)^{p+d}}{p^p d^d} \\ &\leq \frac{\exp\left(\frac{1}{24}\right)}{\sqrt{2\pi}} \sqrt{2} \frac{d^p}{p^p} \left(1 + \frac{p}{d}\right)^{p+d} \\ &\leq \frac{d^p}{p^p} \exp\left(\frac{p^2}{d} + p\right) \end{aligned}$$

which proves the result. \square

Combining the last two lemmas, we get the following bound on the Christoffel function.

Lemma 6.6 *Let $\mathbf{x} \notin S$ and δ be such that $\text{dist}(\mathbf{x}, S) \geq \delta$. Then, for any $d \in \mathbb{N}$, $d > 0$, we have:*

$$s(d) \Lambda_{\mu_S, d}(\mathbf{x}) \leq 2^{3 - \frac{\delta d}{\delta + \text{diam}(S)}} d^p \left(\frac{e}{p}\right)^p \exp\left(\frac{p^2}{d}\right).$$

Proof We may translate the origin of \mathbb{R}^p at \mathbf{x} and scale the coordinates by $\delta + \text{diam}(S)$, this results in $\mathbf{x} = 0$ and distance from \mathbf{x} to S is at most $\delta' = \frac{\delta}{\delta + \text{diam}(S)} \leq 1$. Furthermore, S is contained in the unit Euclidean ball \mathbf{B} . Using invariance of the Christoffel function with respect to change of origin and change of basis in \mathbb{R}^p (Theorem 3.5), this affine transformation does not change the value of the Christoffel function. Now, the polynomial described in Lemma 6.3 provides an upper bound on the Christoffel function. Indeed for any $d' \in \mathbb{N}$, we have:

$$\Lambda_{\mu_S, 2d'+1}(0) \leq \Lambda_{\mu_S, 2d'}(0) \leq 2^{2-2\delta'd'} \leq 2^{3-\delta'(2d'+1)}, \tag{6.7}$$

where we have used $\delta' \leq 1$ to obtain the last inequality. Combining Lemma 6.5 and (6.7), we obtain for any $d' \in \mathbb{N}$:

$$\begin{aligned} s(2d') \Lambda_{\mu_S, 2d'}(0) &\leq 2^{3-2\delta'd'} (2d')^p \left(\frac{e}{p}\right)^p \exp\left(\frac{p^2}{2d'}\right), \\ s(2d'+1) \Lambda_{\mu_S, 2d'+1}(0) &\leq 2^{3-\delta'(2d'+1)} (2d'+1)^p \left(\frac{e}{p}\right)^p \exp\left(\frac{p^2}{2d'+1}\right). \end{aligned} \tag{6.8}$$

Since in (6.8) $d' \in \mathbb{N}$ was arbitrary, we obtain in particular:

$$s(d) \Lambda_{\mu_S, d}(0) \leq 2^{3-\delta'd} d^p \left(\frac{e}{p}\right)^p \exp\left(\frac{p^2}{d}\right). \tag{6.9}$$

The result follows from (6.9) by setting $\delta' = \frac{\delta}{\delta + \text{diam}(S)}$. □

Proof of Theorem 3.9

Proof Let us first prove that $\lim_{k \rightarrow \infty} d_H(S, S_k) = 0$. We take care of both expressions in the definition of d_H separately. Fix an arbitrary $k \in \mathbb{N}$, from Assumption 3.7 and Lemma 6.6, for any $\mathbf{x} \in \mathbb{R}^p$ such that $\text{dist}(\mathbf{x}, S) > \delta_k$:

$$\begin{aligned} s(d_k) \Lambda_{\mu_S, d_k}(\mathbf{x}) &\leq 2^{3 - \frac{\text{dist}(\mathbf{x}, S) d_k}{\text{dist}(\mathbf{x}, S) + \text{diam}(S)}} d_k^p \left(\frac{e}{p}\right)^p \exp\left(\frac{p^2}{d_k}\right) \\ &< \alpha_k. \end{aligned}$$

From this, we deduce that $\mathbb{R}^p \setminus S_k \supseteq \{\mathbf{x} \in \mathbb{R}^p : \text{dist}(\mathbf{x}, S) > \delta_k\}$ and thus $S_k \subseteq \{\mathbf{x} \in \mathbb{R}^p : \text{dist}(\mathbf{x}, S) \leq \delta_k\}$. Since k was arbitrary, for any $k \in \mathbb{N}$:

$$\sup_{\mathbf{x} \in S_k} \text{dist}(\mathbf{x}, S) \leq \delta_k. \tag{6.10}$$

Inequality (6.10) allows taking care of one term in the expression of d_H . Let us now consider the second term. We would like to show that:

$$\sup_{\mathbf{x} \in S} \text{dist}(\mathbf{x}, S_k) \rightarrow 0 \quad \text{as } k \rightarrow \infty. \tag{6.11}$$

Note that the supremum is attained in (6.11). We will prove this by contradiction: for the rest of the proof, M denotes a fixed positive number whose value can change between expressions. Suppose that (6.11) is false. This means that for each $k \in \mathbb{N}$ (up to a subsequence), we can find $\mathbf{x}_k \in S$ which satisfies:

$$\text{dist}(\mathbf{x}_k, S_k) \geq M \tag{6.12}$$

Since $\mathbf{x}_k \in S$ and S is compact, the sequence $(\mathbf{x}_k)_{k \in \mathbb{N}}$ has an accumulation point $\bar{\mathbf{x}} \in S$, i.e., (up to a subsequence) $\mathbf{x}_k \rightarrow \bar{\mathbf{x}}$ as $k \rightarrow \infty$. Since $\text{dist}(\cdot, S_k)$ is a Lipschitz function, combining with (6.12), for every $k \in \mathbb{N}$ (up to a subsequence),

$$\text{dist}(\bar{\mathbf{x}}, S_k) \geq M. \tag{6.13}$$

We next show that (6.13) contradicts the assumption $S = \text{cl}(\text{int}(S))$. From now on, we discard terms not in the subsequence and assume that (6.13) holds for all $k \in \mathbb{N}$. Combining Lemma 6.2 and Assumption 3.7, for every $k \in \mathbb{N}$:

$$S_k \supseteq \{\mathbf{x} \in S : \text{dist}(\mathbf{x}, \partial S) \geq \delta_k\}. \tag{6.14}$$

Since $S = \text{cl}(\text{int}(S))$ and $\bar{\mathbf{x}} \in S$, consider a sequence $\{\mathbf{y}_l\}_{l \in \mathbb{N}} \subset \text{int}(S)$ such that $\mathbf{y}_l \rightarrow \bar{\mathbf{x}}$ as $l \rightarrow \infty$. Since $\mathbf{y}_l \in \text{int}(S)$, we have $\text{dist}(\mathbf{y}_l, \partial S) > 0$ for all l . Up to a rearrangement of the terms, we may assume that $\text{dist}(\mathbf{y}_l, \partial S)$ is decreasing and $\text{dist}(\mathbf{y}_0, \partial S) \geq \delta_0$. For all l , denote by k_l the smallest integer such that $\text{dist}(\mathbf{y}_l, \partial S) \geq \delta_{k_l}$. We must have $k_l \rightarrow \infty$ and we can discard terms so that k_l is a valid subsequence. We have constructed a subsequence k_l such that for every $l \in \mathbb{N}$, $\mathbf{y}_l \in S_{k_l}$ and $\mathbf{y}_l \rightarrow \bar{\mathbf{x}}$. This is in contradiction with (6.13) and hence (6.11) must be true. Combining (6.10) and (6.11), we have that $\lim_{k \rightarrow \infty} d_H(S, S_k) = 0$.

Let us now prove that $\lim_{k \rightarrow \infty} d_H(\partial S, \partial S_k) = 0$, we begin with the term $\sup_{\mathbf{x} \in \partial S_k} \text{dist}(\mathbf{x}, \partial S)$. Fix an arbitrary $k \in \mathbb{N}$ and $\bar{\mathbf{x}} \in \partial S_k$. We will distinguish the cases $\bar{\mathbf{x}} \in S$ and $\bar{\mathbf{x}} \notin S$. Assume first that $\bar{\mathbf{x}} \notin S$. We deduce from (6.10), that:

$$\text{dist}(\bar{\mathbf{x}}, \partial S) = \text{dist}(\bar{\mathbf{x}}, S) \leq \delta_k. \tag{6.15}$$

Assume now that $\bar{\mathbf{x}} \in S$. If $\bar{\mathbf{x}} \in \partial S$, we have $\text{dist}(\bar{\mathbf{x}}, \partial S) = 0$. Assume that $\bar{\mathbf{x}} \in \text{int}(S)$. From (6.14), we have that $S \setminus S_k \subseteq \{\mathbf{x} \in S : \text{dist}(\mathbf{x}, \partial S) < \delta_k\}$ and hence $\text{cl}(S \setminus S_k) \subseteq \{\mathbf{x} \in S : \text{dist}(\mathbf{x}, \partial S) \leq \delta_k\}$. Since $\bar{\mathbf{x}} \in \partial S_k \cap \text{int}(S)$, we have $\bar{\mathbf{x}} \in \text{cl}(S \setminus S_k)$ and hence $\text{dist}(\bar{\mathbf{x}}, \partial S) \leq \delta_k$. Combining the two cases $\bar{\mathbf{x}} \in S$ and $\bar{\mathbf{x}} \notin S$, we have in any case that $\text{dist}(\bar{\mathbf{x}}, \partial S) \leq \delta_k$ and hence:

$$\sup_{\mathbf{x} \in \partial S_k} \text{dist}(\mathbf{x}, \partial S) \leq \delta_k. \tag{6.16}$$

Let us now prove that:

$$\sup_{\mathbf{x} \in \partial S} \text{dist}(\mathbf{x}, \partial S_k) \rightarrow 0 \quad \text{as } k \rightarrow \infty. \tag{6.17}$$

First, since S is closed by assumption, the supremum is attained for each $k \in \mathbb{N}$. Assume that (6.17) does not hold, this means there exists a constant $M > 0$, such that

we can find $\mathbf{x}_k \in \partial S, k \in \mathbb{N}$ with $\text{dist}(\mathbf{x}_k, \partial S_k) \geq M$. If $\mathbf{x}_k \notin S_k$ infinitely often, then, we would have up to a subsequence $\mathbf{x}_k \in S$ and $\text{dist}(\mathbf{x}_k, S_k) \geq M$. This is exactly (6.12) and we already proved that it cannot hold true. Hence, $\mathbf{x}_k \notin S_k$ only finitely many times and we may assume by discarding finitely many terms that $\mathbf{x}_k \in S_k$ for all $k \in \mathbb{N}$. Let $\bar{\mathbf{x}} \in \partial S$ be an accumulation point of $(\mathbf{x}_k)_{k \in \mathbb{N}}$. Since $\bar{\mathbf{x}} \in \partial S$, there exists $\bar{\mathbf{y}} \notin S$ such that $0 < \text{dist}(\bar{\mathbf{y}}, S) \leq \liminf_{k \rightarrow \infty} \text{dist}(\mathbf{x}_k, \partial S_k)/2$. Since $\mathbf{x}_k \in S_k$ for all k sufficiently large, we have $\bar{\mathbf{y}} \in S_k$ for all k sufficiently large but the fact that $0 < \text{dist}(\bar{\mathbf{y}}, S)$ contradicts (6.10). Hence, (6.17) must hold true and the proof is complete. \square

Remark 6.7 (Refinements) The proof of Theorem 3.9 is based on the following fact:

$$\{\mathbf{x} \in \mathbb{R}^p : \text{dist}(\mathbf{x}, \bar{S}) \geq \delta_k\} \subseteq S_k \subseteq \{\mathbf{x} \in \mathbb{R}^p : \text{dist}(\mathbf{x}, S) \leq \delta_k\}.$$

Depending on the regularity of the boundary ∂S of S , it should be possible to get sharper bounds on the distance as a function of δ_k . This should involve the dependency on δ of the function:

$$\delta \rightarrow d_H(\{\mathbf{x} \in \mathbb{R}^p : \text{dist}(\mathbf{x}, \bar{S}) \geq \delta\}, \partial S).$$

For example, if the boundary ∂S has bounded curvature, this function is equal to δ for sufficiently small δ . Another example, if $S \subset \mathbb{R}^2$ is the interior region of a nonself-intersecting continuous polygonal loop, then the function is of the order of $\frac{\delta}{\sin(\frac{\theta}{2})}$, where θ is the smallest angle between two consecutive segments of the loop.

Proof of Theorem 3.12

Proof Lemma 6.2 holds with μ in place of μ_S and w_- in place of $\frac{1}{\lambda(S)}$. Indeed, we have:

$$\Lambda_{\mu,d} \geq w_- \lambda(\mathbf{B}_\delta(\mathbf{x})) \Lambda_{\mu_{\mathbf{B}_\delta(\mathbf{x})}}$$

and the rest of the proof remains the same with different constants. Similarly, Lemma 6.6 holds with μ in place of μ_S ; indeed, the proof only uses the fact that μ_S is a probability measure supported on S which is also true for μ . The proof then is identical to that of Theorem 3.9 by reflecting the corresponding change in the constants. \square

Proof of Theorem 3.13

A preliminary Lemma

Lemma 6.8 *Let μ be a probability measure supported on a compact set S . Then, for every $d \in \mathbb{N}, d > 0$, and every $\mathbf{x} \in \mathbb{R}^p$:*

$$\Lambda_{\mu,d}(\mathbf{x}) \leq \left(\frac{\text{diam}(\text{conv}(S))}{\text{dist}(\mathbf{x}, \text{conv}(S)) + \text{diam}(\text{conv}(S))} \right)^2.$$

Proof Set $\mathbf{y} = \text{proj}(\mathbf{x}, \text{conv}(S))$, that is $\|\mathbf{y} - \mathbf{x}\| = \text{dist}(\mathbf{x}, \text{conv}(S))$ and:

$$\mathbf{y} = \arg \min_{\mathbf{z} \in \text{conv}(S)} \{ \langle \mathbf{z}, \mathbf{y} - \mathbf{x} \rangle \}. \tag{6.18}$$

Consider the affine function:

$$\mathbf{z} \mapsto f_{\mathbf{x}}(\mathbf{z}) := \frac{\left\langle \mathbf{x} - \mathbf{z}, \frac{\mathbf{x} - \mathbf{y}}{\|\mathbf{x} - \mathbf{y}\|} \right\rangle}{\|\mathbf{x} - \mathbf{y}\| + \text{diam}(\text{conv}(S))}. \tag{6.19}$$

For any $\mathbf{z} \in S$, we have:

$$f_{\mathbf{x}}(\mathbf{z}) \leq \frac{\|\mathbf{x} - \mathbf{z}\|}{\|\mathbf{x} - \mathbf{y}\| + \text{diam}(\text{conv}(S))} \leq \frac{\|\mathbf{x} - \mathbf{y}\| + \|\mathbf{y} - \mathbf{z}\|}{\|\mathbf{x} - \mathbf{y}\| + \text{diam}(\text{conv}(S))} \leq 1, \tag{6.20}$$

where we have used Cauchy-Schwartz and triangular inequalities. Furthermore, we have for any $\mathbf{z} \in S$:

$$f_{\mathbf{x}}(\mathbf{z}) \geq \min_{\mathbf{z} \in \text{conv}(S)} f_{\mathbf{x}}(\mathbf{z}) = \frac{\|\mathbf{x} - \mathbf{y}\|}{\|\mathbf{x} - \mathbf{y}\| + \text{diam}(\text{conv}(S))}, \tag{6.21}$$

where we have used equation (6.18). Consider the affine function $q_{\mathbf{x}}: \mathbf{z} \rightarrow 1 - f_{\mathbf{x}}(\mathbf{z})$. We have:

$$q_{\mathbf{x}}(\mathbf{x}) = 1 \tag{6.22}$$

$$0 \leq q_{\mathbf{x}}(\mathbf{z}) \leq \frac{\text{diam}(\text{conv}(S))}{\|\mathbf{x} - \mathbf{y}\| + \text{diam}(\text{conv}(S))}, \text{ for any } \mathbf{z} \in S,$$

where the inequalities are obtained by combining (6.20) and (6.21). The result follows from (6.18), (6.22), and Theorem 3.1. □

Proof of Theorem 3.13

Proof First, let us consider measurability issues. Fix n and d such that $M_d(\mu)$ is invertible. Let \mathbf{X} be a matrix in $\mathbb{R}^{p \times n}$, we use the shorthand notation:

$$\Lambda_{\mathbf{X},d}(\mathbf{z}) = \min_{P \in \mathbb{R}_d[\mathbf{x}], P(\mathbf{z})=1} \frac{1}{n} \sum_{i=1}^n P(\mathbf{X}_i)^2, \tag{6.23}$$

where for each i , \mathbf{X}_i is the i th column of the matrix \mathbf{X} . This corresponds to the empirical Christoffel function with input data given by the columns of \mathbf{X} . Consider the function $F: \mathbb{R}^{p \times n} \rightarrow [0, 1]$ defined as follows:

$$F: \mathbf{X} \rightarrow \sup_{\mathbf{z} \in \mathbb{R}^p} |\Lambda_{\mu,d}(\mathbf{z}) - \Lambda_{\mathbf{X},d}(\mathbf{z})|. \tag{6.24}$$

It turns out that F is a semi-algebraic function (its graph is a semi-algebraic set). Roughly speaking, a set is semi-algebraic if it can be defined by finitely many polynomial inequalities. We refer the reader to [10] for an introduction to semi-algebraic geometry; we mostly rely on content from Chapter 2. First, the function

$$(\mathbf{X}, \mathbf{z}, P) \rightarrow \frac{1}{n} \sum_{i=1}^n P(\mathbf{X}_i)^2$$

is semi-algebraic (by identifying the space of polynomials with the Euclidean space of their coefficients) and the set $\{(P, \mathbf{z}) : P(\mathbf{z}) = 1\}$ is also semi-algebraic. Constrained partial minimization can be expressed by a first-order formula, and, by Tarski-Seidenberg Theorem (see, e.g., [10, Theorem 2.6]), this operation preserves semi-algebraicity. Hence, the function $(\mathbf{X}, \mathbf{z}) \rightarrow \Lambda_{\mathbf{X},d}(\mathbf{z})$ is semi-algebraic. Furthermore, Theorem 3.1 ensures that $\Lambda_{\mu,d}(\mathbf{z}) = 1/\kappa(\mathbf{z}, \mathbf{z})$ for any \mathbf{z} , where $\kappa(\mathbf{z}, \mathbf{z})$ is a polynomial in \mathbf{z} and hence $\mathbf{z} \rightarrow \Lambda_{\mu,d}(\mathbf{z})$ is semi-algebraic. Finally, absolute value is semi-algebraic and using a partial minimization argument again, we have that F is a semi-algebraic function.

As a semi-algebraic function, F is Borel measurable. Indeed, using the *good sets principle* ([2] §1.5.1, p. 35) it is sufficient to prove that for an arbitrary interval³ $(a, b) \subset [0, 1]$, $F^{-1}((a, b)) \in \mathcal{B}(\mathbb{R}^{p \times n})$. Any such set is the pre-image of a semi-algebraic set by a semi-algebraic map. As proved in [10, Corollary 2.9], any such set must be semi-algebraic and hence measurable. Thus, with the notations of Theorem 3.13, $\|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_\infty$ is indeed a random variable for each fixed n, d such that $M_d(\mu)$ is invertible.

We now turn to the proof of the main result of the Theorem. For simplicity, we adopt the following notation for the rest of the proof. For any continuous function $f : \mathbb{R}^p \rightarrow \mathbb{R}$, and any subset $V \subseteq \mathbb{R}^p$:

$$\|f\|_V := \sup_{\mathbf{x} \in V} |f(\mathbf{x})|, \quad [\text{so that } \|f\|_{\mathbb{R}^p} = \|f\|_\infty], \tag{6.25}$$

which could be infinite. We prove that for any $\epsilon > 0$:

$$P(\limsup_n \{\|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{\mathbb{R}^p} \geq \epsilon\}) = 0, \tag{6.26}$$

where the probability is taken with respect to the random choice of the sequence of independent samples from μ and the limit supremum is the set theoretic limit of the underlying events.

Fix $\epsilon > 0$. Denote by S the compact support of μ . Note that S contains also the support of μ_n with probability 1. From Lemma 6.8, we have an upper bound on both $\Lambda_{\mu_n,d}$ and $\Lambda_{\mu,d}$ of order $O(\text{dist}(\mathbf{x}, \text{conv}(S))^{-2})$ which holds with probability 1. Hence, it is possible to find a compact set V_ϵ containing S (with complement $V_\epsilon^c = \mathbb{R}^n \setminus V_\epsilon$) such that, almost surely:

$$\max \{ \|\Lambda_{\mu_n,d}\|_{V_\epsilon^c}, \|\Lambda_{\mu,d}\|_{V_\epsilon^c} \} \leq \frac{\epsilon}{2}. \tag{6.27}$$

Next, we have the following equivalence:

$$\|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{\mathbb{R}^p} \geq \epsilon \iff \begin{cases} \|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{V_\epsilon} \geq \epsilon \text{ or} \\ \|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{V_\epsilon^c} \geq \epsilon \end{cases} \tag{6.28}$$

On the other hand, since both functions are nonnegative, from equation (6.27), almost surely:

$$\|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{V_\epsilon^c} \leq \max \{ \|\Lambda_{\mu_n,d}\|_{V_\epsilon^c}, \|\Lambda_{\mu,d}\|_{V_\epsilon^c} \} \leq \frac{\epsilon}{2}. \tag{6.29}$$

³Recall that the Borel σ -field $\mathcal{B}([0, 1])$ is generated by the intervals (a, b) of $[0, 1]$ (see [2] §1.4.6, p. 27).

Hence, the second event in the right-hand side of (6.28) occurs with probability 0. As a consequence, except for a set of events of measure 0, we have:

$$\|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{\mathbb{R}^p} \geq \epsilon \Leftrightarrow \|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{V_\epsilon} \geq \epsilon,$$

which in turn implies:

$$P \left(\limsup_n \left\{ \|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{\mathbb{R}^p} \right\} \geq \epsilon \right) = P \left(\limsup_n \left\{ \|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{V_\epsilon} \right\} \geq \epsilon \right). \tag{6.30}$$

By assumption, the moment matrix $\mathbf{M}_d(\mu)$ is invertible and by the strong law of large numbers, almost surely, $\mathbf{M}_d(\mu_n)$ must be invertible for sufficiently large n . Assume that $\mathbf{M}_d(\mu_n)$ is invertible, we have:

$$\begin{aligned} \|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{V_\epsilon} &= \sup_{\mathbf{x} \in V_\epsilon} \left\| \frac{1}{\mathbf{v}_d(\mathbf{x})^T \mathbf{M}_d(\mu)^{-1} \mathbf{v}_d(\mathbf{x})} - \frac{1}{\mathbf{v}_d(\mathbf{x})^T \mathbf{M}_d(\mu_n)^{-1} \mathbf{v}_d(\mathbf{x})} \right\| \\ &= \sup_{\mathbf{x} \in V_\epsilon} \left\| \frac{\mathbf{v}_d(\mathbf{x})^T (\mathbf{M}_d(\mu_n)^{-1} - \mathbf{M}_d(\mu)^{-1}) \mathbf{v}_d(\mathbf{x})}{\mathbf{v}_d(\mathbf{x})^T \mathbf{M}_d(\mu)^{-1} \mathbf{v}_d(\mathbf{x}) \mathbf{v}_d(\mathbf{x})^T \mathbf{M}_d(\mu_n)^{-1} \mathbf{v}_d(\mathbf{x})} \right\|. \end{aligned} \tag{6.31}$$

Using the strong law of large numbers again, continuity of eigenvalues and the fact that for large enough n , $\mathbf{M}_d(\mu_n)$ is invertible with probability 1, the continuous mapping theorem ensures that almost surely, for n sufficiently large, the smallest eigenvalue of $\mathbf{M}_d(\mu_n)^{-1}$ is close to that of $\mathbf{M}_d(\mu)^{-1}$ and hence bounded away from 0. Since the first coordinate of $\mathbf{v}_d(\mathbf{x})$ is 1, the denominator in (6.31) is bounded away from 0 almost surely for sufficiently large n . In addition, since V_ϵ is compact, $\mathbf{v}_d(\mathbf{x})$ is bounded on V_ϵ and there exists a constant K such that, almost surely, for sufficiently large n :

$$\|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{V_\epsilon} \leq K \|\mathbf{M}_d(\mu)^{-1} - \mathbf{M}_d(\mu_n)^{-1}\|, \tag{6.32}$$

where the matrix norm in the right-hand side is the operator norm induced by the Euclidean norm. Combining (6.30) and (6.32), we obtain:

$$\begin{aligned} &P \left(\limsup_n \left\{ \|\Lambda_{\mu_n,d} - \Lambda_{\mu,d}\|_{\mathbb{R}^p} \right\} \geq \epsilon \right) \\ &\leq P \left(\limsup_n \left\{ K \|\mathbf{M}_d(\mu)^{-1} - \mathbf{M}_d(\mu_n)^{-1}\| \geq \epsilon \right\} \right). \end{aligned} \tag{6.33}$$

The strong law of large numbers and the continuity of the matrix inverse $\mathbf{M}_d(\cdot)^{-1}$ at μ ensure that the right-hand side of (6.33) is 0. This concludes the proof. \square

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

1. Aaron, C., Bodart, O.: Local convex hull support and boundary estimation. *J. Multivar. Anal.* **147**, 82–101 (2016)
2. Ash, R.B.: *Real Analysis and Probability*. Academic Press Harcourt Brace Jovanovich, Publishers, Boston (1972)
3. Bařfılo, A., Cuevas, A., Justel, A.: Set estimation and nonparametric detection. *Can. J. Stat.* **28**(4), 765–782 (2000)

4. Basu, S., Pollack, R., Roy, M.F.: Computing the first Betti number and the connected components of semi-algebraic sets. In: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing, pp. 304–312 (2005)
5. Berman, R.J.: Bergman kernels for weighted polynomials and weighted equilibrium measures of \mathbb{C}^n . Indiana University Mathematics Journal **58**(4), 1921–1946 (2009)
6. Bos, L.: Asymptotics for the Christoffel function for Jacobi like weights on a ball in $\mathbb{R}\mathbb{R}^m$. New Zealand Journal of Mathematics **23**(99), 109–116 (1994)
7. Bos, L., Della Vecchia, B., Mastroianni, G.: On the asymptotics of Christoffel functions for centrally symmetric weight functions on the ball in $\mathbb{R}\mathbb{R}^d$. Rendiconti del Circolo Matematico di Palermo **2**(52), 277–290 (1998)
8. Chevalier, J.: Estimation du Support et du Contour du Support d'une Loi de probabilité. Annales de l'Institut Henri Poincaré, Section B **12**(4), 339–364 (1976)
9. Cholaquidis, A., Cuevas, A., Fraiman, R.: On poincaré cone property. Ann. Stat. **42**(1), 255–284 (2014)
10. Coste, M.: An introduction to semialgebraic geometry. Istituti Editoriali e Poligrafici Internazionali (2000)
11. Cuevas, A., Fraiman, R.: A plug-in approach to support estimation. Ann. Stat. **25**, 2300–2312 (1997)
12. Cuevas, A., González-Manteiga, W., Rodríguez-casal, A.: Plug-in estimation of general level sets. Aust. N. Z. J. Stat. **48**(1), 7–19 (2006)
13. Davis, J., Goadrich, M.: The relationship between Precision-Recall and ROC curves. In: Proceedings of the 23rd international conference on Machine learning, pp. 233–240. ACM (2006)
14. De Marchi, S., Sommariva, A.: M. Vianello Multivariate Christoffel functions and hyperinterpolation. Dolomites Research Notes on Approximation **7**, 26–3 (2014)
15. Devroye, L., Wise, G.L.: Detection of abnormal behavior via nonparametric estimation of the support. SIAM J. Appl. Math. **38**(3), 480–488 (1980)
16. Dunkl, C.F., Xu, Y.: Orthogonal Polynomials of Several Variables. Cambridge University Press, Cambridge (2001)
17. Geffroy, J.: Sur un problème d'estimation géométrique. Publications de l'Institut de Statistique des Universités de Paris **13**, 191–210 (1964)
18. Gustafsson, B., Putinar, M., Saff, E., Stylianopoulos, N.: Bergman polynomials on an archipelago: estimates, zeros and shape reconstruction. Adv. Math. **222**(4), 1405–1460 (2009)
19. Hardle, W., Park, B.U., Tsybakov, A.B.: Estimation of non-sharp support boundaries. J. Multivar. Anal. **55**(2), 205–218 (1995)
20. Helton, J.W., Lasserre, J.B., Putinar, M.: Measures with zeros in the inverse of their moment matrix. Ann. Probab. **36**(4), 1453–1471 (2008)
21. Kroò, A., Lubinsky, D.S.: Christoffel functions and universality in the bulk for multivariate orthogonal polynomials. Can. J. Math. **65**(3), 600620 (2012)
22. Kroó, A., Lubinsky, D.S.: Christoffel functions and universality on the boundary of the ball. Acta Math. Hungar. **140**, 117–133 (2013)
23. Lasserre, J.B., Pauwels, E.: Sorting out typicality with the inverse moment matrix SOS polynomial. In: Proceedings of the 30-th Conference on Advances in Neural Information Processing Systems (2016)
24. Lichman, M.: UCI Machine Learning Repository, <http://archive.ics.uci.edu/ml> University of California, Irvine, School of Information and Computer Sciences (2013)
25. Malyshkin, V.G.: Multiple Instance Learning: Christoffel Function Approach to Distribution Regression Problem. arXiv:1511.07085 (2015)
26. Mammen, E., Tsybakov, A.B.: Asymptotical minimax recovery of sets with smooth boundaries. Ann. Stat. **23**(2), 502–524 (1995)
27. Máté, A., Nevai, P.: Bernstein's Inequality in L^p for $0 < p < 1$ and $(C, 1)$ Bounds for Orthogonal Polynomials. Ann. Math. **111**(1), 145–154 (1980)
28. Máté, A., Nevai, P., Totik, V.: Szegő's extremum problem on the unit circle. Ann. Math. **134**(2), 433–453 (1991)
29. Molchanov, I.S.: A limit theorem for solutions of inequalities. Scand. J. Stat. **25**(1), 235–242 (1998)
30. Nevai, P.: Géza Freud, orthogonal polynomials and Christoffel functions. A case study. Journal of Approximation Theory **48**(1), 3–167 (1986)
31. Parzen, E.: On estimation of a probability density function and mode. Ann. Math. Stat. **33**(3), 1065–1076 (1962)

32. Patschkowski, T., Rohde, A.: Adaptation to lowest density regions with application to support recovery. *Ann. Stat.* **44**(1), 255–287 (2016)
33. Polonik, W.: Measuring mass concentrations and estimating density contour clusters, an excess mass approach. *Ann. Stat.* **23**(3), 855–881 (1995)
34. Rényi, A., Sulanke, R.: Über die konvexe hülle von n zufällig gewählten Punkten. *Probab. Theory Relat. Fields* **2**(1), 75–84 (1963)
35. Rigollet, P., Vert, R.: Optimal rates for plug-in estimators of density level sets. *Bernoulli* **15**(4), 1154–1178 (2009)
36. Robbins, H.: A remark on Stirling’s formula. *Am. Math. Mon.* **62**(1), 26–29 (1955)
37. Rosenblatt, M.: Remarks on some nonparametric estimates of a density function. *Ann. Math. Stat.* **27**(3), 832–837 (1956)
38. Schölkopf, B., Platt, J., Shawe-Taylor, J., Smola, A., Williamson, R.: Estimating the support of a high-dimensional distribution. *Neural Comput.* **13**(7), 1443–1471 (2001)
39. Singh, A., Scott, C., Nowak, R.: Adaptive Hausdorff estimation of density level sets. *Ann. Stat.* **37**(5B), 2760–2782 (2009)
40. Szegő, G.: Orthogonal polynomials. In: *Colloquium publications*, AMS, (23), fourth edition (1974)
41. Totik, V.: Asymptotics for Christoffel functions for general measures on the real line. *Journal d’Analyse Mathématique* **81**(1), 283–303 (2000)
42. Tsybakov, A.B.: On nonparametric estimation of density level sets. *Ann. Stat.* **25**(3), 948–969 (1997)
43. Williams, G., Baxter, R., He, H., Hawkins, S., Gu, L.: A comparative study of RNN for outlier detection in data mining. In: *IEEE International Conference on Data Mining* (p. 709). IEEE Computer Society (2002)
44. Xu, Y.: Christoffel functions and Fourier series for multivariate orthogonal polynomials. *Journal of Approximation Theory* **82**(2), 205–239 (1995)
45. Xu, Y.: Asymptotics for orthogonal polynomials and Christoffel functions on a ball. *Methods Appl. Anal.* **3**, 257–272 (1996)
46. Xu, Y.: Asymptotics of the christoffel functions on a simplex in $\mathbb{R}R^d$. *Journal of Approximation Theory* **99**(1), 122–133 (1999)
47. Zeileis, A., Hornik, K., Smola, A., Karatzoglou, A.: Kernlab-an S4 package for kernel methods in R. *J. Stat. Softw.* **11**(9), 1–20 (2004)