Reconstruction of Multivariate Functions from Scattered Data

R. Schaback

June 3, 2003

Contents

1 Introduction **5**
1.1 Recovery, Interpolation, and Approximation 6
1.2 Input and Output Data 6
1.3 Restrictions on the Choice of Spaces 7
1.4 Data-dependent Spaces 8
1.5 Evaluation, Interpolation and Approximation 10
1.6 Conditionally Positive Definite Functions 12
1.7 Basic Equations for Conditionally Positive Definite Functions . 13

2 Working with Basis Functions **14**
2.1 General Practical Considerations 15
2.2 Current Basis Functions 18
2.3 Computational Complexity of Solving the System 19
2.4 Multilevel Algorithms 20
2.5 Numerical Examples 21

3 General Theory **25**
3.1 Optimal Recovery in Hilbert Spaces 27
3.2 Spaces of Functions 40
3.3 Native Spaces 52
3.4 Error Bounds in Native Spaces of Functions 63
3.5 Upper Bounds for the Optimal Power Function 71
3.6 Doubling the Approximation Order 93
3.7 Improvement by Localization 94

4 Advanced Results on $\mathbb{R}^d$ **95**
4.1 Fourier Transforms on $\mathbb{R}^d$ 95
4.2 Native Spaces Characterized by Transforms 102
4.3 Condition Numbers 103
4.4 Error Bounds 103

5 Special Theory **103**
5.1 Results for General Transforms 103
5.2 Theory on the Torus using Fourier Series 106
5.3 Theory on Spheres using Expansions 106
5.4 Lower Bounds for Eigenvalues 106
5.5 Generalizations of Results Using Transforms 107

6 Theory on Grids **107**
6.1 Strang-Fix Theory 107
6.2 Application to Radial Basis Functions 107
6.3 Shift Invariant Spaces 107

7 Construction and Characterization of Positive Definite Functions **107**
7.1 General Construction Techniques 107
7.2 Construction of Radial Functions on $\mathbb{R}^d$ 107
7.3 Positive Definite Functions on Topological Groups 107
7.4 Positive Definite Zonal Functions on Spheres 107

8 Special Algorithms **107**
8.1 Reduction of Enlarged System, Method 1 108
8.2 Reduction of Enlarged System, Method 2 109

9 Computational Geometry Techniques **112**
9.1 Voronoi Diagrams 112

10 Appendix **112**
10.1 Basis Functions 112
10.2 MATLAB routines 112
10.3 Hilbert Space Basics 113
10.4 Special Functions and Transforms 122
10.5 Necessary Results from Real Analysis 122

List of Figures

1 Franke's function 23
2 Reconstruction of Franke's function from thin-plate splines on $M = 81$ points 23
3 Residuals for thin-plate splines on $M = 81$ points 23
4 Recovery using Wendland's C^2 function with support radius 2 on $M = 81$ points 23
5 Residuals using Wendland's C^2 function with support radius 2 on $M = 81$ points 23
6 Recovery using Wendland's C^2 function with bandwidth 21 on $M = 289$ points 23
7 Residuals using Wendland's C^2 function with bandwidth 21 on $M = 289$ points 23

8 Residuals using Wendland's C^2 function with bandwidth 21 on $M = 4225$ points 23
9 Recovery using Wendland's C^2 function with bandwidth 45 on $M = 289$ points 23
10 Residuals using Wendland's C^2 function with bandwidth 45 on $M = 4225$ points 23
11 Recovery using Wendland's C^2 function with bandwidth 21 on $M = 289$ points, multilevel method 26
12 Residuals using Wendland's C^2 function with bandwidth 21 on $M = 289$ points, multilevel method 26
13 Recovery using Wendland's C^2 function with bandwidth 5 on $M = 9$ points, multilevel method 26
14 Recovery using Wendland's C^2 function with bandwidth 5 on $M = 25$ points, multilevel method 26
15 Recovery using Wendland's C^2 function with bandwidth 5 on $M = 81$ points, multilevel method 26
16 Recovery using Wendland's C^2 function with bandwidth 5 on $M = 289$ points, multilevel method 26
17 Recovery using Wendland's C^2 function with bandwidth 5 on $M = 4225$ points, multilevel method 26

List of Tables

1 Conditionally Positive Definite Functions (*TCPDFct*) 18
2 Unconditionally Positive Definite Functions (*TPDFct*) 18
3 Errors for interpolation of Franke's function, Nonstationary Case 22
4 Errors for interpolation of Franke's function, Stationary Case . 24
5 Errors for interpolation of Franke's function, Stationary Case, Interpolation of residuals 25
6 L_∞ Bounds of Power Function Based on Lagrange Data (*TCPDEB*) 91
7 Conditionally Positive Definite Functions (*TCPDFct2*) 113
8 Unconditionally Positive Definite Functions (*TPDFct2*) . . . 113

Foreword

This text currently serves two purposes:

- it backs up the lecture on reconstruction of multivariate functions as given in Göttingen in summer 1996, and
- it serves as a gradually growing reference manual for research of the group in Göttingen and related places.

It may finally develop into a monograph, but as to now it is rather preliminary and not intended for general distribution. Suggestions, corrections, addenda, and any form of criticism are welcome.

R. Schaback Göttingen, June 3, 2003

e-mail: schaback@namu01.gwdg.de

Layout for this test version:

- Logical LaTeX labels printed out in *slanted* font.
- Three-level cumulative numbering of environments and equations.
- The index is just a preliminary and (possibly) unsorted list of keywords.

1 Introduction

(*SectIntro*) The following is intended to give the basic motivation for what follows in later chapters. It shows that the reconstruction of multivariate functions f from certain function spaces $\mathcal{F}$ requires dependence of $\mathcal{F}$ on the data. Such data-dependent spaces are provided by conditionally positive definite functions, and these are in the focus of this text. Their optimality properties, as proven in later sections, justify this point of view. After definition of spaces generated by conditionally positive definite functions, this section introduces the standard algorithms for recovery of functions from such spaces. Examples, generalizations, proofs, theoretical details, and implementation problems will be added in later sections.

1.1 Recovery, Interpolation, and Approximation

In almost all practical applications, a function f is given not as a formula, but as a bunch of data. These data often take the form of approximate values $f(x_1), \ldots, f(x_M)$ of f at some scattered locations $x_1, \ldots, x_M$ of the domain Ω of definition of f. The **recovery problem** then consists in the reconstruction of f as a formula from the given data. This reconstruction can be done in two ways:

- **interpolation** tries to match the data exactly, taking f from a large class $\mathcal{F}$ of functions that is actually able to meet all of the data, while

- **approximation** allows f to miss the data somewhat, but selects the reconstruction function from a smaller class $\mathcal{F}$ of functions that will not in general be able to reconstruct the data exactly.

The selection between interpolation and approximation will depend on the application, and especially on the choice of function classes $\mathcal{F}$ and the necessity of exact reproduction of data.

We shall address both problems here, and there will be some hidden links discovered between the two approaches. Furthermore, we shall allow a much wider class of recovery problems in later sections, but the basic motivation is better shown by the above simplified "Lagrange" setting.

1.2 Input and Output Data

We shall consider reconstruction of d-variate functions f defined on a domain Ω. In most cases, Ω will be a subset of $\mathbb{R}^d$, but many results will hold on general sets. Right from the start we keep in mind that d might be large and that the domain Ω may be all of $\mathbb{R}^d$ or something special like a subdomain of $\mathbb{R}^d$ or the $d-1$ dimensional **sphere**, i.e. the surface $\{ x \in \mathbb{R}^d : \|x\|_2 = 1 \}$ of the **unit ball** $\{ x \in \mathbb{R}^d : \|x\|_2 \leq 1 \}$, where $\|.\|_2$ denotes the usual **Euclidean norm** on $\mathbb{R}^d$. In addition, we also may encounter very large sets of data, and these usually come up in two parts:

- a finite set $X = \{x_1, \ldots, x_M\}$ of M possibly wildly scattered points in some domain $\Omega \subseteq \mathbb{R}^d$, and

- real numbers $f_1, \ldots, f_M$ that represent approximate values of f at the given points.

The reconstruction should supply some function f defined on a domain Ω that contains all the data locations, and the data are to be reproduced approximately in the sense

$$f_j \approx f(x_j), \quad 1 \le j \le M.$$

But there are two other important input data for the recovery process:

- the domain Ω should be prescribed by the user, and
- the reconstruction should be confined to some prescribed class $\mathcal{F}$ of functions in order to avoid unpredictable results.

These will finally fix the set of formulas that are allowed as the output of the recovery process. Their choice will very much depend on the application and on additional knowledge of the user. For instance, somebody might want the resulting function f to be defined on all of $\mathbb{R}^d$, while somebody else is interested in a much more local reconstruction, e.g. in the convex hull of the data locations.

Furthermore, there may be different requirements on the smoothness of the recovered function or on its decay further away from the data. These have to be incorporated into the choice of $\mathcal{F}$, in addition to further information the user can provide.

1.3 Restrictions on the Choice of Spaces

There are two good reasons to assume that the class $\mathcal{F}$ of functions should be a linear space:

- If the values f_j are multiplied by a fixed scalar factor α, then the new data should be recovered by the function αf instead of f.
- If data f_j and g_j at the same locations $x_j \in \mathbb{R}^d$ are recovered by functions f and g, respectively, then the data f_j+g_j should be recovered by the function $f + g$.

Note that this does not only require the class $\mathcal{F}$ to be a linear space: it also enforces the whole recovery process to consist of linear maps that associate a function to each data set. Furthermore, the recovery process will have a nonunique solution and thus be numerically unstable, if there is a function g in $\mathcal{F}$ that vanishes at all data locations in $X = \{x_1, \ldots, x_M\}$, because then all functions of the form αg can be added to a solution f without altering the data reproduction.

Definition 1.3.1 *(DefNond) If $\mathcal{F}$ is a space of functions on a domain Ω, then a subset X of Ω is called $\mathcal{F}$-***nondegenerate***, if zero is the only function from $\mathcal{F}$ that vanishes on X.*

We see that only the $\mathcal{F}$-nondegenerate subsets X of Ω can be used for stable reconstruction. It would be nice if any finite set X or at least (if $dim\mathcal{F} = M$) any set $X = \{x_1, \ldots, x_M\}$ would be nondegenerate for a given space $\mathcal{F}$.

But in truly multivariate situations this turns out to be *impossible*. In fact, if a linear subspace $\mathcal{F}$ of dimension $M \geq 2$ of a space of multivariate functions is fixed independent of the set $X = \{x_1, \ldots, x_M\}$, there always is a degenerate set X. This surprising and disappointing observation dates back to Mairhuber and Curtis (cf. [1](*BraessBuch*)):

Theorem 1.3.2 *(MCTheorem) Let $\mathcal{F}$ be an M-dimensional space of continuous real-valued functions on some domain $\Omega \subseteq \mathbb{R}^d$ which is truly d-dimensional in the sense that it contains at least an open subset Ω_1 of $\mathbb{R}^d$. Assume further that any set $X = \{x_1, \ldots, x_M\} \subseteq \Omega_1$ is $\mathcal{F}$-nondegenerate. Then either $M = 1$ or $d = 1$ hold, i.e. either the function space or the underlying domain are just one-dimensional.*

Proof. We can assume $\Omega = \Omega_1$ without loss of generality. If the continuous functions $v_1, \ldots, v_M$ are a basis of $\mathcal{F}$, then the function $D(x_1, \ldots, x_M) = \det(v_j(x_k))$ is a continuous function of its M arguments. Due to our assumption this function can vanish only if two or more of the arguments coincide. Let us assume $M \geq 2$, and let Ω be at least truly 2-dimensional. Then one can swap the points x_1 and x_2 by a continuous motion that avoids coincidence with any of the arguments. Thus there is a continuous transition between $D(x_1, x_2, x_3, \ldots, x_M)$ and $D(x_2, x_1, x_3, \ldots, x_M) = -D(x_1, x_2, x_3, \ldots, x_M)$ that keeps D away from zero. This is impossible. □

1.4 Data-dependent Spaces

(*SubSectDDSpaces*) The Mairhuber-Curtis theorem 1.3.2 (*MCTheorem*) forces us to let the space $\mathcal{F}$ depend on the data. But since for given $X = \{x_1, \ldots, x_M\}$ there should be a linear recovery map

$$R_X : \mathbb{R}^M \to \mathcal{F}. \quad (f_1, \ldots, f_M) \mapsto f,$$

it is reasonable to let $\mathcal{F}$ depend on the data **locations** $X = \{x_1, \ldots, x_M\}$ only, not on the data *values* $f_1, \ldots, f_M$. The formulas for the construction of functions $f(x)$ in $\mathcal{F}$ thus must depend on $X = \{x_1, \ldots, x_M\}$ and generate a

linear space. The most straightforward way to achieve this is to combine the arguments x and x_j into a *single* function

$$\Phi : \Omega \times \Omega \to I\!R$$

and to view each $\Phi(x, x_j)$ as a data-dependent function of the variable x. Superposition of such functions results in a space

(*calfdef*)

$$\mathcal{F}_{X,\Phi} := \left\{ \sum_{j=1}^{M} \alpha_j \Phi(x, x_j) \quad : \quad \alpha_j \in I\!R \right\} \tag{1.4.1}$$

that may serve our purposes. It will turn out later that there are strong arguments to support this definition of a data-dependent space of functions. Under quite weak and general assumptions it can be proven that there is no better way to do it. Details of this will be given in 3.1.4 (*Necessity)*, but we cite the basic features here to support some useful simplifications. If for some Φ the union of all function spaces $\mathcal{F}_{X,\Phi}$ for varying sets X is required to have **translation invariance**, then the function Φ should be of the special form

$$\Phi(x, y) = \phi(x - y), \quad \phi : I\!R^d \to I\!R^d.$$

If we add **rotational invariance**, we end up with **radial basis functions**

$$\Phi(x, y) = \phi(\|x - y\|_2), \quad \phi : I\!R_{\geq 0} \to I\!R.$$

Note that in the latter case there is only a single *univariate* function required to generate a large class of spaces of *multivariate* functions. If we are working on the unit sphere in $I\!R^d$ and assume rotational invariance, we get **zonal** functions

$$\Phi(x, y) = \phi(x^T y), \quad \phi : I\!R_{\geq 0} \to I\!R^d,$$

where x^T stands for transposition of the vector x such that $\Phi(x, y)$ just is a univariate function $\phi(x^T y)$ of the scalar product $x^T y$. Details are provided in section 3.2.4 (*SecIP)*.

Of course there are other methods to generate data-dependent linear spaces of functions. The most prominent one is used widely in the theory of **finite elements**. There, the data set $X = \{x_1, \ldots, x_M\}$ is first used to generate a triangulation of its convex hull, and then one constructs functions on each subset of the triangulation, which are finally patched together to form smooth global functions. This approach is very effective if the space dimension d is small and if the functions to be recovered need not be very smooth. We refer the reader to the vast literature on this subject, and we proceed without considering triangulations of domains and patching of functions.

1.5 Evaluation, Interpolation and Approximation

(*subsecEIA*) The representation of functions in (1.4.1, *calfdef*) now serves as the reconstruction formula, and all one has to do when solving the reconstruction problem is to determine the vector $\alpha = (\alpha_1, \ldots, \alpha_M)$ of the coefficients of the resulting function with the representation

(*falphadef*)

$$f_\alpha(x) := \sum_{j=1}^{M} \alpha_j \Phi(x, x_j), \quad x \in \Omega \subseteq I\!R^d. \tag{1.5.1}$$

Before we turn to this problem, we note that evaluation of such a function at large numbers of different locations $x \in \Omega$ can be quite cumbersome if M is large. However, the strong dependence on M can be relaxed if the values $\Phi(x, x_j)$ vanish whenever x and x_j are not near to each other. Examples of such functions will be given later.

Reconstruction by *interpolation* on $X = \{x_1, \ldots, x_M\}$ will now require to solve the linear system

(*EQsys1*)

$$\sum_{j=1}^{M} \alpha_j \Phi(x_k, x_j) = f_k, \quad k = 1, \ldots, M \tag{1.5.2}$$

for $\alpha_1, \ldots, \alpha_M$. We shall write this in shorthand matrix form as

$$A\alpha \;=\; f,$$

but in cases where the dependence on X and Φ is crucial, we add capital subscripts:

$$A_{X,\Phi}\alpha_{X,\Phi} \;=\; f_X, \quad A_{X,\Phi} = (\Phi(x_k, x_j))_{1 \le j,k \le M}\,.$$

To make the system uniquely solvable, the matrix A must be nonsingular. Looking at approximation, we shall soon have additional reasons to assume that $A_{X,\Phi}$ should even be positive definite. Thus it is more or less unavoidable to assume $A_{X,\Phi}$ to be positive definite for all X, when the function Φ is fixed. For these reasons we require the function Φ to satisfy

Definition 1.5.3 *(DPD) A real-valued function*

$$\Phi : \Omega \times \Omega \to I\!R$$

is a **positive definite function** *on* Ω*, iff for any choice of finite subsets* $X = \{x_1, \ldots, x_M\} \subseteq \Omega$ *of* M *different points the matrix*

$$A_{X,\Phi} = (\Phi(x_k, x_j))_{1 \le j,k \le M}$$

is positive definite.

At first sight it seems to be a miracle that a fixed function Φ should be sufficient to make all matrices of the above form positive definite, no matter which points are chosen and no matter how many. It is even more astonishing that one can often pick radial functions like $\Phi(x, y) = \exp(\|x - y\|_2^2)$ to do the job, and to work for **any** space dimension.

Turning to approximation, the space $\mathcal{F}_{X,\Phi}$ of (1.4.1, *calfdef*) should depend on less data than those given to determine the approximation. We simply assume some other data on some (large) Lebesgue-measurable subset $\Omega_1 \subseteq \Omega$ to be specified, and approximation should take place in the space $L_2(\Omega_1)$, for instance, which we assume to contain $\mathcal{F}_{X,\Phi}$. This covers discrete and continuous least-squares fits on the set Ω_1 by functions of the form f_α from (1.5.1, *falphadef*). The normal equations for the approximation are

$$\sum_{j=1}^{M} \alpha_j (\Phi(\cdot, x_k), \Phi(\cdot, x_j))_{L_2(\Omega_1)} = (\Phi(\cdot, x_k), f(\cdot))_{L_2(\Omega_1)}, \quad k = 1, \ldots, M.$$

Introducing new functions

(*Psidef*)

$$\Psi(x, y) := (\Phi(\cdot, x), \Phi(\cdot, y))_{L_2(\Omega_1)} \tag{1.5.4}$$

$$g(y) := (\Phi(\cdot, y), f(\cdot))_{L_2(\Omega_1)}$$

we see that this is exactly an *interpolation* system of the form

$$A_{X,\Psi} \alpha_{X,\Psi} = g_X.$$

Thus approximation reduces to interpolation by functions from a similar, but somewhat different function space.

At this point we see how positive definiteness comes in: the above matrix $A_{X,\Psi}$ is a Gramian with respect to the functions $\Phi(\cdot, x_k)$ in the inner-product space $L_2(\Omega_1)$. Thus it is positive definite whenever these functions are linearly independent in $L_2(\Omega_1)$. But the latter requirement is unavoidable for stable approximation in $L_2(\Omega_1)$.

From these preliminary considerations we conclude that positive definite functions should be investigated further, and we note in passing that (1.5.4, *Psidef*) yields a first method to construct such functions Ψ from linear independent functions $\Phi(\cdot, x_k),\ 1 \le k \le M, x_k \in \Omega$. We shall consider such constructions in detail in section 7.1 (*SecGCT*), but we remark in passing that the **Gaussian**

$$\Phi(x, y) := \exp(-\alpha \|x - y\|_2^2)$$

is positive definite on $\mathbb{R}^d$ for all d and all $\alpha > 0$. Since the proof requires tools like Fourier transforms, we defer it to section 7.2.1 (*SecPDG*).

1.6 Conditionally Positive Definite Functions

Positive definite functions (formerly defined in a slightly different way) have a long history that is nicely surveyed by Stewart [29](*Stewart:76-1*). However, the first cases of radial basis functions used widely and successfully in applications were

- the **thin-plate spline** $\Phi(x, y) = \phi(\|x - y\|_2) = -\|x - y\|_2 \log \|x - y\|_2$ introduced by Duchon [6](*duchon:76-1*), [7](*duchon:78-1*), [8](*duchon:79-1*),
- the **multiquadric** $\Phi(x, y) = \phi(\|x - y\|_2) = \sqrt{c^2 + \|x - y\|_2^2}$ and
- the **inverse multiquadric** $\Phi(x, y) = \phi(\|x - y\|_2) = \dfrac{1}{\sqrt{c^2 + \|x - y\|_2^2}}$ used by the geophysicist Hardy [13](*hardy:71-1*)

but the first two of these are *not* positive definite. The corresponding matrices $A_{X,\Phi}$ naturally define quadratic forms

(*QFdef*)

$$Q_{X,\Phi} \ : \ (\alpha_1, \ldots, \alpha_M) \mapsto \alpha^T A_{X,\Phi} \alpha := \sum_{j,k=1}^{M} \alpha_j \alpha_k \Phi(x_j, x_k) \tag{1.6.1}$$

on $\mathbb{R}^d$, where ^{T} stands for vector transposition, but these forms are positive definite only on a proper subspace of $\mathbb{R}^M$. More precisely, for certain positive values of m the above functions Φ satisfy the following

Definition 1.6.2 *(DCPD) A real-valued function*

$$\Phi : \Omega \times \Omega \to \mathbb{R}$$

is a **conditionally positive definite function** *of order m on $\Omega \subseteq \mathbb{R}^d$, iff for any choice of finite subsets $X = \{x_1, \ldots, x_M\} \subseteq \Omega$ of M different points the value*

$$\alpha^T A_{X,\Phi} \alpha := \sum_{j,k=1}^{M} \alpha_j \alpha_k \Phi(x_j, x_k)$$

of the quadratic form (1.6.1, QFdef) is positive, provided that the vector $\alpha = (\alpha_1, \ldots, \alpha_M) \in \mathbb{R}^M \setminus \{0\}$ has the additional property

(CPDef)

$$\sum_{j=1}^{M} \alpha_j p(x_j) = 0 \tag{1.6.3}$$

for all d-variate polynomials p of order (=degree-1) up to m. The linear space of such polynomials will be denoted by $\mathbb{P}_m^d$, and its dimension is $q := \binom{m-1+d}{d}$.

It is a major problem to prove that multiquadrics are conditionally positive definite of a fixed order m for all space dimensions d. This was done (among other things) in Micchelli's fundamental paper [19](*micchelli:86-1*) that boosted the research on radial basis functions.

1.7 Basic Equations for Conditionally Positive Definite Functions

If Φ is conditionally positive definite of order m on $\Omega \subseteq \mathbb{R}^d$, then the additional condition (1.6.3, *CPDef*) reduces the M degrees of freedom of $\alpha \in \mathbb{R}^M$ by at most q, the dimension of the space $\mathbb{P}_m^d$ of polynomials. Thus it is reasonable to add q new degrees of freedom to the recovery process by adding $\mathbb{P}_m^d$ to the space of admissible functions. Then (1.4.1, *calfdef*) has to be replaced by

(*calfdef2*)

$$\mathcal{G}_{X,\Phi} := \mathbb{P}_m^d + \mathcal{F}_{X,\Phi} = \mathbb{P}_m^d + \left\{ \sum_{j=1}^{M} \alpha_j \Phi(x, x_j) \quad : \quad \alpha_j \in \mathbb{R} \text{ with (1.6.3, } \textit{CPDef}) \right\}. \tag{1.7.1}$$

Now the $M \times M$ system (1.5.2, *EQsys1*) goes over into the $(M+q) \times (M+q)$ system

(*EQsys2*)

$$\begin{array}{rcrcll} \sum_{j=1}^{M}\alpha_j\Phi(x_k,x_j) & + & \sum_{i=1}^{q}\beta_i p_i(x_k) & = & f_k, & 1\le k\le M \\ \sum_{j=1}^{M}\alpha_j p_i(x_j) & + & 0 & = & 0, & 1\le i\le q \end{array} \tag{1.7.2}$$

for vectors $\alpha = (\alpha_1,\ldots,\alpha_M) \in \mathbb{R}^M$ and $\beta = (\beta_1,\ldots,\beta_q) \in \mathbb{R}^q$, where the polynomials $p_1,\ldots,p_q$ are a basis of $\mathbb{P}_m^d$. Introducing a matrix

$$P := P_X := (p_i(x_j))_{1\le i\le q, 1\le j\le M},$$

of values of polynomials, this system reads in matrix form as

(*BDef*)

$$\begin{pmatrix} A & P \\ P^T & 0 \end{pmatrix}\begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}. \tag{1.7.3}$$

The coefficient matrix of this enlarged linear system will be abbreviated by B or $B_{X,\Phi}$. The solvability of (1.7.2, *EQsys2*) is described by

Theorem 1.7.4 *(Nonsing1) Let Φ be conditionally positive definite of order m on $\Omega \subseteq \mathbb{R}^d$, and let the data set $X = \{x_1,\ldots,x_M\} \subseteq \Omega$ be $\mathbb{P}_m^d$-nondegenerate. Then the system (1.7.2, EQsys2) is uniquely solvable. Furthermore, there are linear algebra techniques using at most $\mathcal{O}(Mq^2 + M^2q)$ operations to reduce it to a positive definite $(M-q)\times(M-q)$ system.*

Proof. Let a pair of vectors $\alpha \in \mathbb{R}^M$ and $\beta \in \mathbb{R}^q$ solve the homogeneous system with matrix (1.7.3, *BDef*). Then we have $A\alpha + P\beta = 0$ and $P^T\alpha = 0$. Multiplying the first equation with α^T and inserting the second in transposed form, we get $\alpha^T A\alpha + 0 = 0$. Now $\alpha = 0$ follows from conditional positive definiteness, and we are left with $P\beta = 0$. This in turn implies $\beta = 0$, because X is $\mathbb{P}_m^d$-nondegenerate. The second assertion will be proven by two explicit algorithms in 8.1 (*Red1)* and 8.2 (*Red2)*. □

2 Working with Basis Functions

This section is intended for readers working on applications. It contains tables of the currently known conditionally positive definite functions and provides guidelines for picking the right function Φ from the tables. These guidelines are based on both numerical experience and theoretical insight.

However, this chapter will not attempt to prove any of the statements inherent in the guidelines, but rather refer the reader to subsequent sections.

Right after giving the general guidelines, we turn to algorithms. Some extensions to the basic algorithms are presented together with efficiency considerations. Special strategies for system reduction, iterative solution, sparse matrices, and preprocessing techniques for large sets of data points are included.

A series of examples serves for illustration. Since these examples are quite convincing in general, they justify the considerable amount of theoretical background to be developed in the later sections.

2.1 General Practical Considerations

Before picking a suitable function Φ for recovering a function f in an application, the user first has to consider the following issues:

- How smooth should f be?
- What is the required behaviour near the boundary of the convex hull or outside of the data set $X = \{x_1, \ldots, x_M\}$?
- Are the data locations evenly or very unevenly distributed?
- Is exact reproduction of the data required?
- Are M and/or the space dimension d so large that efficiency considerations are predominant over reproduction quality questions?

2.1.1 Uncertainty Relation

(*GPCUP*) When considering the above questions, the user has to keep in mind that every good thing has its price. This basic fact of real life occurs here in the form of an *Uncertainty Relation*:

> If you go for good reproduction quality, you have to sacrifice numerical stability. If you go for good stability, you have to sacrifice reproduction quality.

This wishy-washy statement will be made precise in 3.4.6 (*URT*), and there it turns out that both reproduction quality and numerical instability are linked to both data density and smoothness of Φ (and, in cases with compact support, to the size of the support radius of Φ). Furthermore, if large linear

systems with positive definite coefficient matrices are solved by the conjugate gradient method, numerical stability is directly linked to efficiency via the condition of the matrices. This is why for large problems one can replace "stability" by "efficiency" in the Uncertainty Relation.

If the data density is considered fixed, the Uncertainty Relation suggests that the user should be very careful about the smoothness of the function Φ. It should be as low as the application tolerates, and any excessive smoothness will have negative effects on stability.

But if reproduction quality or stability is fixed, there is a trade-off between data density and smoothness of Φ. For sparse data one can work with smooth functions, and for large and dense data sets one has to work with low smoothness of Φ in order to avoid numerical problems. If working with compactly supported functions Φ, this is a standard way to escape the inherent numerical problems with very large and dense data sets. One can split the data set into subsets of increasing density and use compactly supported functions with decreasing support radii on these data sets. If things work out nice, one can expect to work at a fixed stability level, but with incereasing local resolution. We treat such *multilevel techniques* in detail in 2.4 *(MLA)* but the next paragraph will add some other arguments in favor of it.

Compactly supported functions offer computational advantages due to sparsity of the corresponding matrices. If supports are small, the effect of such functions will be strictly local, and this has both advantages and disadvantages. The disadvantage is that global effects cannot be nicely recovered, and thus small supports should be used only in cases where the global behavior is already recovered by any other method. The usual trick is to

- first apply a global method (possibly using a small but global data set),
- take the residuals (data minus values of the recovery function) and then
- handle the local effects by reconstruction the residuals using compactly supported functions on the full data set.

This three-stage process is quite common in applications and amounts to solve for the global *trend* first and then to model the local effects on a finer scale. The last two steps can be iterated using smaller and smaller supports, and this is the multilevel method that we look at in 2.4 *(MLA)*

2.1.2 Unevenly Distributed Data

(*GPCUDD*) The above statements assume a more or less evenly scattered data set. If there are local clusters of data points or areas without data, some other aspects come into the game. In fact, for a fixed function Φ the numerical stability and the reproduction quality are connected to two similar, but different quantities which roughly coincide for evenly distributed data sets. The stability is connected to the **separation distance**

(*SDDef*)

$$s := s_X := \frac{1}{2} \min_{1 \le j \ne k \le M} \|x_j - x_k\|_2 \tag{2.1.1}$$

while the reproduction quality on the domain Ω is ruled by a somewhat more complicated quantity (see (3.5.8, *hrhodef*)) that can roughly be described for practical purposes by the **fill distance**

(*DDDef*)

$$h := h_{X,\Omega} := \sup_{x \in \Omega} \min_{1 \le j \le M} \|x - x_j\|_2. \tag{2.1.2}$$

Separation distance measures the minimal distance that separates any two data locations, i.e. it is the minimal distance from any point of the data set to its nearest data point, while fill distance measures the way how the data fill the domain, i.e. it is the maximal distance from any point of the domain to its nearest data point. Thus fill distance is never smaller than separation distance, but hazardous cases have a very small separation distance relative to the fill distance. We call a data set *unevenly distributed* if this happens, and the quotient

$$\delta_{X,\Omega} := \frac{h_{X,\Omega}}{s_X} \ge 1$$

is a good measure for the unevenness of a data distribution X with respect to a domain Ω.

Now the naive treatment of unevenly distributed data sets will induce "additional" numerical instabilities caused by the irregularity of the data distribution. If these instabilities are severe, some action must be taken. If caused by a few points that are extremely near to other data locations with comparable data values, the user can simply throw these "duplicates" out of the data set and proceed, expecting that the nearby data points are sufficient for the required reconstruction.

But there are cases where the data show local clusters which themselves consist of nicely distributed data locations. Then the problem lives on more

than one density scale, and the obvious technique to treat such cases is by working in several steps with increasing local resolution. This is another good reason for the multilevel approach in 2.4 (*MLA)*.

2.2 Current Basis Functions

(*SecCBF*) Table 1 (*TCPDFct)* lists some of the currently known radial functions that are conditionally positive definite of *positive* order m on $\mathbb{R}^d$. A more or less complete list will be in the Appendix under 10.1 (*SecBF)*. Note that these have some polynomial growth towards infinity, and they always generate non-sparse matrices. They work for any space dimension d, and they are especially useful for cases where decay towards infinity is a disadvantage. Thus they should not be applied to residuals but rather to the original data, and their power lies in good reproduction of the global overall shape of the function to be reconstructed, especially in areas away from the data locations.

We now turn to *unconditionally* positive definite functions defined on $\mathbb{R}^d$.

$\phi(r)$	Parameters	m
r^β	$\beta > 0,\ \beta \notin 2\mathbb{N}$	$m \geq \lceil \beta/2 \rceil$
$r^\beta \log r$	$\beta > 0,\ \beta \in 2\mathbb{N}$	$m > \beta/2$
$(r^2 + c^2)^{\beta/2}$	$\beta > 0,\ \beta \notin 2\mathbb{N}$	$m \geq \lceil \beta/2 \rceil$

Table 1: Conditionally Positive Definite Functions (*TCPDFct*)

$\phi(r)$	Parameters	Smoothness	Dimension	Name/Reference
$e^{-\beta r^2}$	$\beta > 0$	$C^\infty(\mathbb{R}^d)$	$d < \infty$	Gaussian
$(r^2 + c^2)^{\beta/2}$	$\beta < 0$	$C^\infty(\mathbb{R}^d)$	$d < \infty$	inv. Multiquadric
$r^\nu K_\nu(r)$	$\nu > 0$	$C^{\lfloor \nu \rfloor}$	$d < \infty$	Sobolev spline
$(1-r)_+^2(2+r)$		C^0	$d \leq 3$	Wu [**?**](*wu:94-1*)
$(1-r)_+^4(1+4r)$		C^2	$d \leq 3$	Wendland [30](*wendland:95-1*)

Table 2: Unconditionally Positive Definite Functions (*TPDFct*)

These have decay towards infinity and come in two variations: compactly supported or not. Due to results given in 7.2.2 (*NECSAlld)* there are no compactly supported positive definite functions that work for all space dimensions. Thus one has to check the space dimension d when working with

compactly supported functions. Table 2 (*TPDFct)* lists some of the currently known cases and provides information about smoothness and admissible space dimensions for positive definiteness. See 10.1 (*SecBF)* for further cases and details.

The decay towards infinity may be an unwanted feature when applied to raw data, but it is very convenient when applied to residuals. Compact supports provide sparse matrices, but the adjustment of the support radius can be hazardous. If chosen too small, the resulting matrices $A_{X,\Phi}$ tend to be nicely diagonal, making the numerical solution very stable and efficient, but the reproduction quality is disastrous, because one reproduces the data by extremely narrow and isolated "delta" peaks. On the contrary, a large support radius very much improves reproduction quality, but at the expense of matrix fill-in and increasing condition. This is another consequence of the Uncertainty Relation.

2.3 Computational Complexity of Solving the System

(*CompEffort*) We now investigate the numerical effort required to solve the system (1.7.3, *BDef*). Assuming that q usually is zero or small compared to M, we roughly have a positive definite and symmetric $M \times M$ system to solve. If the condition is reasonable and M is not too large, Cholesky factorization will do the job at about $M^3/6 + \mathcal{O}(M^2)$ computational cost. However, this is not acceptable for large M. In particular, the value of M can be even too large to form the full matrix in storage. Therefore one has to look for iterative methods and sparse matrix techniques. Some special tricks due to Beatson [3](*beatson-newsam:92-1*) and Powell [26](*powell:92-1*)[25](*powell:92-2*)[6](*PowellEffTPSSystem*) are possible for specific basis functions, but we concentrate here on the solution via compactly supported functions.

In this case the matrix is sparse and its bandwidth depends on the relative size δ/h of the support radius δ and the fill distance h. For a fixed compactly supported positive definite function Φ the effect of an increase of δ yields

- an increase of the bandwidth of the matrix in (1.7.3, *BDef*) via an increase of δ/h,
- an increase of the reproduction quality via an increase of δ/h (see 3.4 (*SecError)*), and
- an increase of its condition via an increase of δ/q (see 3.4.4 (*SecCondition)*).

This is another version of the Uncertainty Relation, and the user has to fix the support radius δ to be sufficiently large to get good reproduction quality while keeping it small enough to let the solution of the system be computationally effective. A general rule of thumb is to work at the limits of the computational resources, and to switch to multilevel techniques (see 2.4 (*MLA)*) in cases where the reproduction quality still is inadequate.

If the ratio δ/q is kept bounded, the norm of the inverse (and thus a major part of the condition) of the matrix in (1.7.3, *BDef*) is bounded. Solving the system by conjugate gradients uses only a fixed number of iterations for fixed precision requirements, if the condition is bounded. Furthermore, each iteration takes only $\mathcal{O}(M \cdot B)$ operations for bandwidth B. Thus the numerical cost cof solving the system (1.7.3, *BDef*) an be kept roughly at $\mathcal{O}(M)$, if the user keeps the ratios of h, q, and δ within reasonable bounds.

We finally check the complexity of evaluating (1.5.1, *falphadef*) at a single argument x. In general one has to expect $\mathcal{O}(M)$, but since one has to evaluate the function in at least $\mathcal{O}(M)$ or many more points, the cost for evaluation will even be underestimated by $\mathcal{O}(M^2)$. For large values of M this cannot be tolerated. Using stencils [26](*powell:92-1*) and Laurent expansions [25](*powell:92-2*) Powell has overcome these difficulities in case of thin-plate splines. For compactly supported basis functions with maximally B points in their support (this coincides with the bandwidth of the system (1.7.3, *BDef*)) one has $\mathcal{O}(B)$ operations for each evaluation, which is a significant advantage if many evaluations have to be made. However, each evaluation then requires to solve the B-nearest-neighbor problem of computational geometry, because for each x one has to pick the B data points x_j with nonzero $\Phi(x, x_j)$ in an effective way. If the data are not too wildly scattered, one can employ preprocessing techniques of complexity at most $\mathcal{O}(M)$ to solve this problem at $\mathcal{O}(1)$ for each x. In general, preprocessing of cost $\mathcal{O}(M \log M)$ is necessary to provide a $\mathcal{O}(\log M)$ complexity of solving the B-nearest-neighbor problem for each x. Details will be provided in section 9 (*SecCGT)*

2.4 Multilevel Algorithms

(*MLA*) The basic idea here is to work at levels indexed by j, where one uses a basis function Φ_j that usually will be compactly supported with a support radius δ_j. On level j the data is confined to a subset X_j of the full data set X, and the corresponding fill distance and separation distance will be denoted by h_j and q_j, respectively. The function f_j to be recovered by some

other function s_j at level j consists of the *residuals* of the preceding step, i.e.

$$f_j := f_{j-1} - s_{j-1},\ j \geq 1,\ f_0 := f.$$

The *ratios* of the three quantities h_j, q_j, and δ_j are kept at reasonable values that make the computations possible, while the quantities themselves decrease with increasing j.

The rationale behind this multilevel techniques is to recover the function f at different levels of resolution, starting from global reconstruction of slowly varying features from coarse global data and ending up with highly local reconstruction of fine details from densely distributed data. The numerical performance of this technique is superior to single-level techniques in applications with very large data sets (see 2.5 (*SecExamples)* and [9](*floater-iske:95-1*) [10](*floater-iske:96-1*)), but its theoretical treatment, starting in [7](*NRSW*), still is incomplete. The numerical cost can be kept to $\mathcal{O}(M)$ by proper choice of supports and fill distances.

2.5 Numerical Examples

(*SecExamples*) Here are some first examples of reconstructions of functions from multivariate scattered data. For easy presentation, we restrict ourselves to two-dimensional cases and use MATLAB for the computations. The corresponding MATLAB M-files and MEX-files are in the appendix.

We start with the reconstruction of Franke's function [11](*franke:82-1*) rescaled to $\Omega = [0,1] \times [0,1] \subset I\!R^2$ from data on a grid $(i/2n, j/2n),\ 0 \leq i,j \leq 2n$ such that $M = (2n+1)^2$. The matrix in (1.7.3, *BDef*) then has approximately $4n^4$ entries, and the computational cost of Cholesky factorization is about $4n^6/3$. If the matrix is non-sparse, only very moderate values of n can be treated.

The function itself (Figure 1 (*FigFranke33)*) is nicely reconstructed up to graphical precision by thin-plate splines $\phi(r) = r^2 \log r$ from information on $M = 81$ data points (Figure 2 (*FigTPS81Fct33)*). The effects of higher values of M are visualized by plotting *residuals* (see Figure 3 (*FigTPS81Res)* for $M = 81$, and note the scale on the z-axis for plots of resuduals).

Working on more than $M = 225$ points becomes very ineffective for non-sparse cases. Thus we now consider examples with Wendland's compactly supported radial basis function $\phi(r) = (1-r)_+^4(4r+1)$ with support scaled to radius δ. On $M = 81$ data points one can still compare with the previous case

while using a large support radius $\delta = 2$ (Figures 4 (*FigWF81Supp2Fct33)*, 5 (*FigWF81Res2)*). To handle larger values of M, the support radius has to be decreased to introduce sparsity. We start with examples having bandwidth 21 on $M = 289$ and $M = 4225$ points (Figures 6 (*FigWF289Ban21Fct65)*, 7 (*FigWF289Ban21)*, 8 (*FigWF4225Ban21)*). Note that the reproduced function is overlaid by some high-frequency wiggles that arise from the small support of the radial basis functio used: the approximation is somewhat too spiky. A look at the residuals supports this, but also implies that the larger errors occur at the boundary. These take over when going to 4225 data points, and make the errors in the interior hardly visible. This is the first hint that the behavior near the boundary needs special treatment.

Now Figures 9 (*FigWF289Ban45Fct65)*, 10 (*FigWF4225Ban45)*) show residuals computed with matrices of bandwidth 45. The results are better, of course, but the message is the same.

For even larger values of M we refrain from providing plots of residuals. Instead, we evaluate the error on a fine grid. Table 3 (*TabNonstat)* on page 22 shows the maximum errors for cases with fixed support radius δ

$N \setminus \delta$	0.03125	0.0625	0.125	0.25	0.5	1	2	4	8
9	*	*	*	*	12.1754	5.4808	5.5436	5.8102	5.9030
25	*	*	*	10.2176	1.1995	0.8186	0.6902	0.6889	0.7073
81	*	*	11.5563	1.1013	0.4668	0.3621	0.3570	0.3584	0.3587
289	*	11.7369	0.8148	0.4606	0.1175	0.0397	0.0241	0.0224	0.0226
1089	11.6653	0.7812	0.4783	0.1158	-	-	-	-	-
4225	0.7791	0.4561	-	-	-	-	-	-	-
16641	-	-	-	-	-	-	-	-	-
66049	-	-	-	-	-	-	-	-	-

Table 3: Errors for interpolation of Franke's function, Nonstationary Case (*TabNonstat*)

* Errors too large due to extremely small supports used,

\- Workspace exhausted or non-sparse matrix.

(*nonstationary case*), as far as the computations were numerically feasible. Convergence along columns is clearly visible, but the scope is still severely limited by computational restrictions.

If the support radius is kept strictly proportional to the fill distance (this is called the fully *stationary* case), then the bandwidth B is constant along

Figure 1: Franke's function
(*FigFranke33*)

Figure 2: Reconstruction of Franke's function from thin-plate splines on $M = 81$ points
(*FigTPS81Fct33*)

Figure 3: Residuals for thin-plate splines on $M = 81$ points
(*FigTPS81Res*)

Figure 4: Recovery using Wendland's C^2 function with support radius 2 on $M = 81$ points
(*FigWF81Supp2Fct33*)

Figure 5: Residuals using Wendland's C^2 function with support radius 2 on $M = 81$ points
(*FigWF81Res2*)

Figure 6: Recovery using Wendland's C^2 function with bandwidth 21 on $M = 289$ points
(*FigWF289Ban21Fct65*)

Figure 7: Residuals using Wendland's C^2 function with bandwidth 21 on $M = 289$ points
(*FigWF289Ban21*)

Figure 8: Residuals using Wendland's C^2 function with bandwidth 21 on $M = 4225$ points
(*FigWF4225Ban21*)

Figure 9: Recovery using Wendland's C^2 function with bandwidth 45 on $M = 289$ points
(*FigWF289Ban45Fct65*)

Figure 10: Residuals using Wendland's C^2 function with bandwidth 45 on $M = 4225$ points
(*FigWF4225Ban45*)

columns in Table 4 (*TabStat)* on page 24, but there is no convergence along

$N \setminus B$	1	5	9	13	21	25	29	37	45
9	12.1754	8.1801	5.4801	5.3389	5.3521	5.3770	5.4083	5.4830	5.5436
25	10.2176	4.6070	1.1993	0.9549	0.9209	0.8995	0.8719	0.8400	0.8186
81	11.5563	4.8475	1.1003	0.8840	0.7236	0.6820	0.6316	0.5368	0.4668
289	11.7369	4.5695	0.8148	0.7554	0.7670	0.7190	0.6606	0.5457	0.4606
1089	11.6653	4.4424	0.7812	0.7831	0.7924	0.7432	0.6838	0.5661	0.4783
4225	11.7024	4.4322	0.7791	0.7733	0.7566	0.7099	0.6529	0.5416	0.4561
16641	11.7109	4.4292	0.7786	0.7119	0.7577	0.6994	0.6578	0.5461	-
66049	12.9205	4.4283	-	-	-	-	-	-	-

Table 4: Errors for interpolation of Franke's function, Stationary Case (*TabStat*)

N number of data points

B number of points per support

\- Workspace exhausted

columns, while the scope is greatly enlarged. Convergence occurs along lines with negative slope in this table, but the minimum attainable error still is quite large. The condition is roughly constant in each column, such that the overall numerical cost is approximately proportional to M.

We now recalculate the columns of Table 4 (*TabStat)* by taking successive residuals as we proceed along each column, working at fixed bandwidth and fixed condition, thus with $\mathcal{O}(M)$ overall computational complexity (see Table 5 (*TabMulti)* on page 25). This multilevel approach now decreases the error significantly and seems to have at least a linear convergence along columns. More information on the numerical behavior of the multilevel approach can be found in [9](*floater-iske:95-1*) [10](*floater-iske:96-1*). Here, we support the results of Table 5 (*TabMulti)* by some additional plots of multilevel interpolants to Franke's function. Figure 11 (*FigWF289Ban21MLFig)* shows the multilevel reconstruction with bandwidth 21 after four levels with 9, 25, 81, and 289 data points. The residuals are in Figure 12 (*FigWF289Ban21ML)* and should be compared with Figure 7 (*FigWF289Ban21)* with the same bandwidth on 289 points, using a single step.

To visualize the smoothing effect of the multilevel method, we pick a drastic example by choosing a very small bandwidth of 5. The reader will realize

$N \setminus B$	1	5	9	13	21	25	29	37	45
9	12.0412	8.1801	5.4801	5.3389	5.3521	5.3770	5.4048	5.4830	5.5436
25	7.6972	2.5971	0.9328	0.7840	0.7016	0.6943	0.6842	0.6808	0.6i845
81	5.9089	0.9172	0.4223	0.3820	0.3565	0.3571	0.3595	0.3680	0.3735
289	4.4449	0.2927	0.0680	0.0518	0.0352	0.0332	0.0314	0.0303	0.0288
1089	3.3053	0.0867	0.0256	0.0187	0.0120	0.0112	0.0105	0.0098	0.0092
4225	2.4589	0.0320	0.0090	0.0064	0.0039	0.0036	0.0034	0.0031	0.0029
16641	1.7481	0.0118	0.0034	0.0023	0.0013	0.0011	0.0011	0.0009	0.0008
66049	1.3085	0.0053	-	-	-	-	-	-	-

Table 5: Errors for interpolation of Franke's function, Stationary Case, Interpolation of residuals
(*TabMulti*)

N number of data points

B number of points per support

\- Workspace exhausted

that this method will be feasible even for gigantic data sets. Figures 13 (*FigWF9Ban5MLFig)* 14 (*FigWF25Ban5MLFig)* 15 (*FigWF81Ban5MLFig)* 16 (*FigWF289Ban5MLFig)* 17 (*FigWF4225Ban5MLFig)* show reconstruction from $M = 9,\ 25,\ 81,\ 289$, and 4225 points. The extremely small bandwidth of 5 does not have a serious influence on the quality on a 3×3 data set, but the spiky reproduction in the medium range introduces wiggles that are ironed out by increasing data density.

Of course, one should take larger supports in the intermediate range and use a bandwidth larger that 5 to produce optimal results, but the above sequence is picked to illustrate what happens qualitatively if the computational restrictions force to work with very small bandwidth. The actual errors can be read off the second column of Table 5 (*TabMulti).*

To prove statements about the convergence rate and the condition of such calculations will be main goal of this text.

3 General Theory

(*SecGT*) Here we start with the basic theoretical foundations and proceed top-down. First, we pose the problem of recovery of elements of Hilbert

Figure 11: Recovery using Wendland's C^2 function with bandwidth 21 on $M = 289$ points, multilevel method
(*FigWF289Ban21MLFig*)

Figure 12: Residuals using Wendland's C^2 function with bandwidth 21 on $M = 289$ points, multilevel method
(*FigWF289Ban21ML*)

Figure 13: Recovery using Wendland's C^2 function with bandwidth 5 on $M = 9$ points, multilevel method
(*FigWF9Ban5MLFig*)

Figure 14: Recovery using Wendland's C^2 function with bandwidth 5 on $M = 25$ points, multilevel method
(*FigWF25Ban5MLFig*)

Figure 15: Recovery using Wendland's C^2 function with bandwidth 5 on $M = 81$ points, multilevel method
(*FigWF81Ban5MLFig*)

Figure 16: Recovery using Wendland's C^2 function with bandwidth 5 on $M = 289$ points, multilevel method
(*FigWF289Ban5MLFig*)

Figure 17: Recovery using Wendland's C^2 function with bandwidth 5 on $M = 4225$ points, multilevel method
(*FigWF4225Ban5MLFig*)

spaces in a very general sense. It turns out that optimal recovery is necessarily linked to the use of conditionally positive definite functions. Conversely, each conditionally positive definite function allows to define a "native" Hilbert space in which it serves to solve an optimal recovery problem. We study the error and the condition of the recovery process and prove the Uncertainty Relation in general. Altogether, this section is intended to contain all theoretical results that can be proven without resort to (Fourier) transforms and which hold for general domains. This implies that the more sophisticated results for special cases are found in later sections.

3.1 Optimal Recovery in Hilbert Spaces

3.1.1 Optimal Recovery Problems

(*subsecORP*) Assume that we want to reconstruct a function f defined on some domain Ω from M pieces of information concerning f. These may for instance be function values $f(x_j),\ 1 \le j \le M$ in case of classical Lagrange interpolation, or inner products $(f, p_j)_{L_2},\ 1 \le j \le M$ for L_2 approximation. In both cases the information consists of the value of a linear functional λ_j applied to f, and in the second case the function f is assumed to lie in a space with an inner product $(\cdot,\cdot)$ that serves to give a specific representation $\lambda_j(f) = (f, p_j)$ to the functionals in question.

To incorporate the second case, we thus assume that there is a space $\mathcal{F}$ of functions and a space $\mathcal{L}$ of functionals such that $\lambda(f)$ is the application of the functional $\lambda \in \mathcal{L}$ to the function $f \in \mathcal{F}$. The space $\mathcal{F}$ is supposed to carry an inner product $(\cdot,\cdot)_{\mathcal{F}}$, and the functionals $\lambda \in \mathcal{L}$ are supposed to be continuous with respect to this inner product, i.e.,

$$|\lambda(f)| \le \|\lambda\|_{\mathcal{L}} \|f\|_{\mathcal{F}}$$

for all $\lambda \in \mathcal{L},\ f \in \mathcal{F}$, where the norm of functionals is defined as usual:

$$\|\lambda\|_{\mathcal{L}} := \sup_{\|f\|_{\mathcal{F}} \neq 0} \frac{|\lambda(f)|}{\|f\|_{\mathcal{F}}} < \infty.$$

We now assume that we want to recover an element f from the space $\mathcal{F}$ using the M real values

(*fj*)

$$\gamma_j = \lambda_j(f),\ 1 \le j \le M \tag{3.1.1}$$

of M linear functionals $\lambda_1, \ldots, \lambda_M$ that are continuous on $\mathcal{F}$. Furthermore, we assume the linear functionals $\lambda_1, \ldots, \lambda_M$ to be linearly independent in $\mathcal{L}$, which means that the information is not redundant.

Then there will usually be many elements $f \in \mathcal{F}$ that satisfy the equations (3.1.1, *fj*), which may now be viewed as generalized interpolation conditions. If f solves (3.1.1, *fj*) and if there is some element $v \in \mathcal{F}$ that satisfies the homogeneous conditions

$$0 = \lambda_j(v),\ 1 \le j \le M,$$

than all elements $f_\alpha := f + \alpha v$ for arbitrary $\alpha \in I\!R$ will solve (3.1.1, *fj*), too. These elements can have arbitrarily large norms, if v is not identically zero. To exclude solutions with extremely large norms one thus asks for elements $f^* \in \mathcal{F}$ that solve (3.1.1, *fj*) and minimize the norm $\|\cdot\|_{\mathcal{F}}$ under all other solutions. That is, the element f^* solves the **optimal recovery problem**

(*ORPF*)

$$\|f^*\|_{\mathcal{F}} = \min_{\substack{f \in \mathcal{F} \\ f_j = \lambda_j(f)}} \|f\|_{\mathcal{F}} \tag{3.1.2}$$

in the space $\mathcal{F}$.

If we pursue this general setting further, we shall finally see that under mild additional assumptions there is a positive definite function that serves to solve the optimal recovery problem. But then we have lost the conditionally positive definite functions of positive order. Thus we try a fresh start that slightly generalizes the above recovery problem.

Instead of a space $\mathcal{F}$ with an inner product, we only assume there is a linear space $\mathcal{G}$ over $I\!R$ with a positive semidefinite bilinear form

$$(\cdot,\cdot)_{\mathcal{G}} : \mathcal{G} \times \mathcal{G} \to I\!R.$$

Then $|g|_{\mathcal{G}}^2 = (g,g)_{\mathcal{G}}$ defines a seminorm $|\cdot|_{\mathcal{G}}$ on $\mathcal{G}$, and we assume that the nullspace

$$\mathcal{P} := \{\ g \in \mathcal{G}\ :\ |g|_{\mathcal{G}} = 0\ \}$$

has a finite dimension $q \ge 0$ and is spanned by a basis $p_1, \ldots, p_q$. As in (3.1.1, *fj*) we assume that we want to recover an element g from the space $\mathcal{G}$ using the M real values

(*gj*)

$$\gamma_j = \lambda_j(g),\ 1 \le j \le M \tag{3.1.3}$$

of M linear functionals $\lambda_1, \ldots, \lambda_M$. But we would run into problems if we would simply assume continuity of these functionals with respect to the seminorm, because this would restrict us to functionals that vanish on $\mathcal{P}$. Postponing the precise assumptions on the functionals, we can now pose the **generalized optimal recovery problem**

(*ORP*)

$$|g^*|_{\mathcal{G}} = \min_{\substack{g \in \mathcal{G} \\ g_j = \lambda_j(g)}} |g|_{\mathcal{G}} \tag{3.1.4}$$

in the space $\mathcal{G}$.

3.1.2 Projection onto the Nullspace

(*SecHSP*) To discuss the solvability of the optimal recovery problem 3.1.4 (*ORP)* in a very general way, we need some more information on the space $\mathcal{G}$ and ist finite-dimensional subspace $\mathcal{P}$. It simplifies later arguments to have a simple way of projecting an element $g \in \mathcal{G}$ onto an element of $\mathcal{P}$. In standard applications, this projection will be an interpolation or an approximation by a low-order polynomial. Such a linear projector $\Pi_{\mathcal{P}}$ from $\mathcal{G}$ onto $\mathcal{P}$ can be defined in many different ways. Here we simply assume that there are q linear functionals $\pi_1, \ldots, \pi_q$ on $\mathcal{G}$ that are linearly independent over $\mathcal{P}$, i.e. the $q \times q$ matrix P with entries $\pi_k(p_j)$ is nonsingular. Then the projector can be represented as

(*DefPN*)

$$\Pi_{\mathcal{P}}(g) := \sum_{j=1}^{q} \pi_j(g) p_j. \tag{3.1.5}$$

By a change of basis in either the p_j or the π_j one can assume that the linear functionals $\pi_j(g)$ satisfy the system

$$\sum_{j=1}^{q} \pi_j(g)\pi_k(p_j) = \pi_k(g),\ 1 \le k \le q.$$

This is just another way of saying

$$\pi_k(\Pi_{\mathcal{P}}(g)) = \pi_k(g),\ 1 \le k \le q,\ g \in \mathcal{G},$$

and it has the consequence that $\Pi_{\mathcal{P}}(p) = p$ for all $p \in \mathcal{P}$, because of $\pi_j(p_k) = \delta_{jk}$.

Having $\Pi_{\mathcal{P}}$ at hand, we now form $R_{\mathcal{P}}(g) := g - \Pi_{\mathcal{P}}(g)$ for all $g \in \mathcal{G}$. For subsequent use we note that the bilinear form on $\mathcal{G}$ can now be rewritten as

(*Rsp*)

$$(f, g)_{\mathcal{G}} = (R_{\mathcal{P}}(f), R_{\mathcal{P}}(g))_{\mathcal{G}}, \quad f, g \in \mathcal{G}. \tag{3.1.6}$$

The decomposition of an arbitrary element $g \in \mathcal{G}$ as

$$g = \Pi_{\mathcal{P}}(g) + R_{\mathcal{P}}(g)$$

implies that the decomposition

(*gdec*)

$$\mathcal{G} = \mathcal{P} + R_{\mathcal{P}}(\mathcal{G}) \tag{3.1.7}$$

is a direct sum, since $R_{\mathcal{P}}(g) \in \mathcal{P}$ implies $g = \Pi_{\mathcal{P}}(g) + R_{\mathcal{P}}(g) \in \mathcal{P}$ and thus $R_{\mathcal{P}}(g) = 0$. Furthermore, the bilinear form $(\cdot, \cdot)_{\mathcal{F}}$ now is positive definite on $R_{\mathcal{P}}(\mathcal{G})$.

3.1.3 Hilbert Space Completion

(*SecHSC*) We now complete the space $R_{\mathcal{P}}(\mathcal{G})$ in the usual way to form a Hilbert space $\mathcal{F}$, taking us back to the setting that we started from, and where

$$(\cdot, \cdot)_{\mathcal{F}} := (R_{\mathcal{P}}(\cdot), R_{\mathcal{P}}(\cdot))_{\mathcal{G}}$$

is the inner product. This completion works via Cauchy sequences modulo null sequences, and it allows all continuous mappings on $R_{\mathcal{P}}(\mathcal{G})$ to be extended to the completion. See Theorem 10.3.11 (*HSCT)* for details. We now define the closure of $\mathcal{G}$ as the direct sum of $\mathcal{P}$ with the closure $\mathcal{F}$ of $R_{\mathcal{P}}(\mathcal{G})$. Then the decomposition (3.1.7, *gdec*) extends to the closures, and if we denote the closure of $\mathcal{G}$ by $\mathcal{G}$ again, we get

(*GPF1*)

$$\mathcal{G} = \mathcal{P} + \mathcal{F}. \tag{3.1.8}$$

Thus we finally see that it makes no difference to start right away with a space $\mathcal{G}$ that allows a decomposition (3.1.7, *gdec*) such that (3.1.6, *Rsp*) is a scalar product on the Hilbert space $\mathcal{F} := R_{\mathcal{P}}(\mathcal{G})$ that has $\mathcal{P}$ as its nullspace.

We finish this section by checking the proper form of admissible functionals for recovery. If λ is just any functional on $\mathcal{G}$, it defines a functional $\lambda - \lambda\Pi_{\mathcal{P}} = \lambda R_{\mathcal{P}}$ by

(*lrest*)

$$g \mapsto \lambda(g) - \lambda(\Pi_{\mathcal{P}}(g)) = \lambda(R_{\mathcal{P}}(g)), \quad g \in \mathcal{G}, \tag{3.1.9}$$

and this functional is a good candidate for being continuous with respect to the seminorm $|\cdot|_{\mathcal{G}}$, because it vanishes on $\mathcal{P}$. We thus consider all functionals λ on $\mathcal{G}$ such that $\lambda - \lambda\Pi_{\mathcal{P}}$ is continuous, and we denote the space of these functionals by $\mathcal{G}^*$. By (3.1.9, *lrest*), for each $\lambda \in \mathcal{G}^*$ the functional $\lambda - \lambda \circ \Pi_{\mathcal{P}}$ is continuous on the Hilbert space $\mathcal{F} = R_{\mathcal{P}}(\mathcal{G})$, and by the Riesz theorem 10.3.14 (*RieszT)* there is an element $g_\lambda \in \mathcal{G}$ such that the identity

(*lrep*)

$$\lambda(g) - \lambda(\Pi_{\mathcal{P}}(g)) = \lambda(R_{\mathcal{P}}(g)) = (g, g_\lambda)_{\mathcal{G}} \tag{3.1.10}$$

holds for all $\lambda \in \mathcal{G}^*$ and all $g \in \mathcal{G}$. We shall use this identity in the more convenient form

$$\lambda(g) = \lambda(\Pi_{\mathcal{P}}(g)) + (g, g_\lambda)_{\mathcal{G}}$$

and note that g_λ is uniquely defined modulo $\mathcal{P}$, while $R_{\mathcal{P}}(g_\lambda)$ is unique. The functionals from (3.1.9, *lrest*) vanish on $\mathcal{P}$ and they form the dual $\mathcal{F}^*$ of $\mathcal{F}$. If one defines $\Pi^*_{\mathcal{P}}(\lambda) := \lambda \circ \Pi_{\mathcal{P}}$ and $\mathcal{P}^* = \Pi^*_{\mathcal{P}}(\mathcal{G}^*)$, then there are decompositions

$$\lambda = \Pi^*_{\mathcal{P}}(\lambda) + (\cdot, g_\lambda)_{\mathcal{G}}$$

$$\mathcal{G}^* = \mathcal{P}^* + \mathcal{F}^*$$

that correspond to those of $g \in \mathcal{G}$ and $\mathcal{G}$ itself.

3.1.4 Solutions of Optimal Recovery Problems

(*Necessity*) We now can return to the problem (3.1.4, *ORP*) of optimal recovery. The given functionals λ_j are assumed to be in $\mathcal{G}^*$. Then they satisfy (3.1.10, *lrep*) and introduce elements $g_j := g_{\lambda_j} \in \mathcal{G}$, $1 \le j \le M$ in the sense

(*lrepj*)

$$\lambda_j(g) - \lambda_j(\Pi_{\mathcal{P}}(g)) = \lambda_j(R_{\mathcal{P}}(g)) = (g, g_j)_{\mathcal{G}}, \ g \in \mathcal{G}. \tag{3.1.11}$$

These elements are not unique, and we could make them unique by defining $g_j := R_{\mathcal{P}}(g_{\lambda_j})$, $1 \le j \le M$, but the following results do not require this uniqueness. We now can characterize the solutions of the recovery problem:

Theorem 3.1.12 *(ORT1) Any solution g^* of the optimal recovery problem (3.1.4, ORP) with functionals $\lambda_1, \dots, \lambda_M \in \mathcal{G}^*$ satisfying (3.1.11, lrepj) has the form*

(grep)

$$g^* = \sum_{j=1}^{M} \alpha_j g_j + \sum_{i=1}^{q} \beta_i p_i \tag{3.1.13}$$

where the coefficients satisfy the linear system

(EQsys3)

$$\begin{array}{rcrcl} \sum_{j=1}^{M} \alpha_j (g_k, g_j)_{\mathcal{G}} & + & \sum_{i=1}^{q} \beta_i \lambda_k(p_i) & = & \gamma_k, \quad 1 \le k \le M \\ \sum_{j=1}^{M} \alpha_j \lambda_j(p_i) & + & 0 & = & 0, \qquad 1 \le i \le q. \end{array} \tag{3.1.14}$$

and any solution of the above system solves the optimal recovery problem.

Note how similar (3.1.14, *EQsys3*) and (1.7.2, *EQsys2*) are, and note that we postpone the discussion of the solvability of (3.1.14, *EQsys3*).

Proof: We start by noting that g^* is a solution of (3.1.4, *ORP*) if and only if it satisfies the variational equation

(*charmin*)

$$(g^*, v)_{\mathcal{G}} = 0 \text{ for all } v \in \mathcal{G} \text{ with } \lambda_j(v) = 0,\ 1 \le j \le M. \tag{3.1.15}$$

This follows from Corollary 10.3.7 (*BAC*) in section 10.3 (*SecHSB*).

If $g^* \in \mathcal{G}$ satisfies (3.1.14, *EQsys3*) and $v \in \mathcal{G}$ satisfies the homogeneous conditions $\lambda_j(v) = 0,\ 1 \le j \le M$, then

$$\begin{array}{rcl} (g^*, v)_{\mathcal{G}} & = & \sum_{j=1}^{M} \alpha_j (g_j, v)_{\mathcal{G}} + \sum_{i=1}^{q} \beta_i (p_i, v)_{\mathcal{G}} \\ & = & \sum_{j=1}^{M} \alpha_j \left(\lambda_j(v) - \lambda_j(\Pi_{\mathcal{P}}(v)) \right) \\ & = & -\sum_{j=1}^{M} \alpha_j \lambda_j(\Pi_{\mathcal{P}}(v)) \\ & = & 0 \end{array}$$

and g^* satisfies (3.1.15, *charmin*) and solves (3.1.4, *ORP*).

To prove the converse, we note that (3.1.15, *charmin*) implies the existence of $\alpha_1, \ldots, \alpha_M \in I\!R$ such that

(*charmin2*)

$$(g^*, v)_{\mathcal{G}} = \sum_{j=1}^{M} \alpha_j \lambda_j(v) \tag{3.1.16}$$

for all $v \in \mathcal{G}$. In fact, the linear map $v \mapsto (g^*, v)_{\mathcal{G}}$ vanishes on the kernel of the linear map $v \mapsto (\lambda_1(v), \ldots, \lambda_M(v))^T \in I\!R^M$ with finite-dimensional range and thus factorizes over the range of this mapping. See the proof of Corollary 10.3.7 (*BAC)* for this argument. But now (3.1.16, *charmin2*) implies

(*charmin3*)

$$(g^*, v)_{\mathcal{G}} = \sum_{j=1}^{M} \alpha_j \left(\lambda_j(\Pi_{\mathcal{P}}(v)) + (g_j, v)_{\mathcal{G}}\right) \tag{3.1.17}$$

and specialization to $v \in \mathcal{P}$ implies the second set of equations in (3.1.14, *EQsys3*). Then (3.1.17, *charmin3*) can be rewritten in the form

$$\left(g^* - \sum_{j=1}^{M} \alpha_j g_j, v\right)_{\mathcal{G}} = 0 \text{ for all } v \in \mathcal{G}$$

and this implies the representation (3.1.13, *grep*) of g^*. The interpolation conditions finally furnish the first set of equations in (3.1.14, *EQsys3*). □

The system (3.1.14, *EQsys3*) looks terrible at first sight, because neither the functions g_j nor their inner products $(g_j, g_k)_{\mathcal{G}}$ are readily available from the given functionals λ_j. But we shall see in (3.2.14, *gjkrep*) that there is a conditionally positive definite function Φ such that

$$(g_j, g_k)_{\mathcal{G}} = \lambda_j^x \lambda_k^y \Phi(x, y)$$

holds for the elements of the matrix in (3.1.14, *EQsys3*), making an easy access to these elements possible, if Φ is explicitly known. In particular, if $\lambda_j(f) = f(x_j)$, then

$$(g_j, g_k)_{\mathcal{G}} = \Phi(x_j, x_k)$$

as we used in (1.7.2, *EQsys2*) in a slightly more special situation.

We now look at solvability of the system (3.1.14, *EQsys3*) in the shorthand form

(*BDef2*)

$$\begin{pmatrix} A & P \\ P^T & 0 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} \gamma \\ 0 \end{pmatrix} \tag{3.1.18}$$

generalizing (1.7.3, *BDef*). If vectors $\alpha \in I\!R^M$ and $\beta \in I\!R^q$ satisfy the homogeneous system, then

$$\begin{aligned} A\alpha &+ P\beta &= 0 \\ P^T\alpha &+ 0 &= 0 \end{aligned}$$

imply

$$\begin{aligned} \alpha^T A\alpha &= 0 \\ P^T\alpha &= 0. \end{aligned}$$

Since the matrix A is a Gramian for the elements $g_1, \ldots, g_M$, it is positive semidefinite and we have

$$\alpha^T A\alpha = \left| \sum_{j=1}^M \alpha_j g_j \right|_{\mathcal{P}}^2 = 0.$$

Thus the element $\sum_{j=1}^M \alpha_j g_j$ of $\mathcal{G}$ must be in $\mathcal{P}$ and the linear combination $\sum_{j=1}^M \alpha_j \lambda_j$ of functionals is zero due to $P^T\alpha = 0$ and

$$\sum_{j=1}^M \alpha_j \lambda_j(v) = \sum_{j=1}^M \alpha_j(\lambda_j(\Pi_{\mathcal{P}}(v)) + (g_j, v)_{\mathcal{G}}) = 0 + \left(\sum_{j=1}^M \alpha_j g_j, v \right)_{\mathcal{G}} = 0$$

for all $v \in \mathcal{G}$. But we assumed the linear functionals $\lambda_1, \ldots, \lambda_M$ to be linearly independent over $\mathcal{G}$. This implies $\alpha = 0$ and we are left with $P\beta = 0$. There is no way to deduce $\beta = 0$ from this in general, and consequently we have to add injectivity of P to our hypotheses, if we want to assure unique solvability of (3.1.14, *EQsys3*). We summarize:

Theorem 3.1.19 *(ORT2) There is a unique solution to the optimal recovery problem (3.1.4, ORP) if the $M \times q$ matrix P with entries*

$$\lambda_j(p_i),\ 1 \le j \le M,\ 1 \le i \le q$$

is injective. This condition means that the only element $p \in \mathcal{P}$ with vanishing data $\lambda_1(p), \ldots, \lambda_M(p)$ must be the zero element. □

It should be clear by now that we finally want to show how the system (1.7.2, *EQsys2*) is a special case of (3.1.14, *EQsys3*) and how a conditionally positive definite function Φ can arise in the above Hilbert space setting. We shall take point evaluation functionals $\lambda_x(v) := (v - \Pi_{\mathcal{P}}(v))(x)$ if the abstract elements $v \in \mathcal{G}$ can be interpreted as functions on some domain Ω containing the points x, and use the elements $g_x := g_{\lambda_x} \in \mathcal{G}$ from (3.1.10, *lrep*) to define a

generalized conditionally positive definite function with $\mathcal{P}$ generalizing $I\!P_m^d$ by

$$\Phi(x,y) := (g_x, g_y)_{\mathcal{G}},\ x,y \in \Omega.$$

The details will be specified in Theorem 3.2.17 (*CPDNeccT)*.

Theorems 3.1.12 (*ORT1)* and 3.1.19 (*ORT2)* show that optimal recovery in the fairly general sense of (3.1.4, *ORP*) necessarily leads to solutions of the special form (3.1.13, *grep*) and linear systems (3.1.14, *EQsys3*). This is why the techniques of section 1.4 (*SubSectDDSpaces)* are a quite natural and general way to access recovery problems.

3.1.5 Related Problems

(*SecRP*) There is an equivalent dual reformulation of the above recovery problem. Instead of reconstructing some $g \in \mathcal{G}$ from the information $\gamma_j = \lambda_j(g),\ 1 \le j \le M$ one can ask for a functional $\lambda^* \in \mathcal{G}^*$ of minimal seminorm in $\mathcal{G}^*$ that satisfies the equations

$$\lambda^*(g_j) = \gamma_j,\ 1 \le j \le M$$

for a set of linearly independent elements $g_1, \dots, g_M \in \mathcal{G}$. For this the dual bilinear form on functionals in $\mathcal{G}^*$ can be defined as

(*DefDualBil*)

$$(\lambda,\mu)_{\mathcal{G}^*} := (g_\lambda, g_\mu)_{\mathcal{G}} = \lambda(g_\mu) - \lambda(\Pi_{\mathcal{P}}(g_\mu)) = \mu(g_\lambda) - \mu(\Pi_{\mathcal{P}}(g_\lambda)). \quad (3.1.20)$$

The additional property required for uniqueness now is that the $M \times q$ matrix P with entries

$$\pi_i(g_j),\ 1 \le j \le M,\ 1 \le i \le q$$

is injective. This condition means that the zero is the only element in the span of $g_1, \dots, g_M$ that projects via $\Pi_{\mathcal{P}}$ to zero in $\mathcal{P}$. We leave details to the readers as an exercise. There is a full duality if one replaces λ_j by g_j and π_i by p_i

Another equivalent optimal recovery problem consists in finding an element $g^* \in \mathcal{G}$ with minimal seminorm $|g^*|_{\mathcal{G}}$ such that

(*scaleq*)

$$\begin{aligned} (g^*, g_j)_{\mathcal{G}} &= \gamma_j,\ 1 \le j \le M \\ \Pi_{\mathcal{P}}(g^*) &= 0, \end{aligned} \quad (3.1.21)$$

where we again assume that the functions g_j represent linear independent functionals λ_j in the sense of (3.1.11, *lrepj*). The difference is that the data now are not taking notice of additional functions from $\mathcal{P}$, such that the second condition of (3.1.21, *scaleq*) is necessary to remove the nonuniqueness of g^* modulo $\mathcal{P}$. Furthermore, one can assume

(*picond*)

$$\Pi_{\mathcal{P}}(g_j) = 0,\ 1 \le j \le M \tag{3.1.22}$$

without loss of generality.

Theorem 3.1.23 *(ORT3) Under the additional assumptions*

(spancond)

$$\sum_{j=1}^{M} \alpha_j g_j \in \mathcal{P} \text{ implies } \alpha_j = 0,\ 1 \le j \le M \tag{3.1.24}$$

and (3.1.22, picond), the above optimal recovery problem with conditions (3.1.21, scaleq) has a unique solution g^* of the form

(grep2)

$$g^* = \sum_{j=1}^{M} \alpha_j g_j \tag{3.1.25}$$

where the coefficients satisfy the linear system

(EQsys4)

$$\sum_{j=1}^{M} \alpha_j (g_k, g_j)_{\mathcal{G}} \ = \ \gamma_k,\ 1 \le k \le M. \tag{3.1.26}$$

Proof: The equivalent variational equation here is

$$(g^*, v)_{\mathcal{G}} = 0 \text{ for all } v \in \mathcal{G} \text{ with } \Pi_{\mathcal{P}}(v) = 0 \text{ and } (v, g_j)_{\mathcal{G}} = 0,\ 1 \le j \le M.$$

This transforms into

$$(g^*, v)_{\mathcal{G}} = (\sum_{j=1}^{M} \alpha_j g_j, v)_{\mathcal{G}}$$

for all $v \in \mathcal{G}$. This is satisfied if (3.1.25, *grep2*) holds. To prove the converse, we conclude that the variational equation implies that the difference of both sides in (3.1.25, *grep2*) lies in $\mathcal{P}$. But application of $\Pi_{\mathcal{P}}$ turns the difference into zero, proving necessity of (3.1.25, *grep2*).

To prove nonsingularity of the system (3.1.26, *EQsys4*) we proceed similarly as in the proof of Theorem 3.1.19 (*ORT2)*, but use (3.1.24, *spancond*) instead of linear independence of the functionals λ_j. □

Note that (3.1.24, *spancond*) is more restrictive than to assume linear independence of the functionals λ_j, as required for Theorem 3.1.12 (*ORT1)*. This is why Theorem 3.1.23 (*ORT3)* has positive definiteness of the matrix $((g_i, g_j)_{\mathcal{G}})_{i,j}$, while Theorems 3.1.12 (*ORT1)* and 3.1.19 (*ORT2)* need the enlarged matrix. Furthermore, the functionals $\mu_j := (\cdot, g_j)_{\mathcal{G}}$ that implicitly arise in Theorem 3.1.23 (*ORT3)* have the additional property $\mu_j(\mathcal{P}) = \{0\}$, and this property is not shared by the functionals λ_j in the previous theorems. In case of $\mathcal{P} = \{0\}$ there is no difference at all.

We now consider the best approximation problem

(*BAP*)

$$\inf_{\lambda \in \Lambda} |\mu - \lambda|_{\mathcal{G}^*} \tag{3.1.27}$$

for a given functional $\mu \in \mathcal{G}^*$ by functionals in

(*DefL*)

$$\Lambda := \operatorname{span}\{\lambda_1, \ldots, \lambda_M\} \subset \mathcal{G}^*. \tag{3.1.28}$$

The usual theory of approximation in spaces with inner products or bilinear forms yields the normal equations

$$(\mu, \lambda_j)_{\mathcal{G}^*} = \sum_{k=1}^{M} \alpha_k(\mu)(\lambda_k, \lambda_j)_{\mathcal{G}^*} = \sum_{k=1}^{M} \alpha_k(\mu)(g_k, g_j)_{\mathcal{G}}$$

with a coefficient matrix as in (3.1.26, *EQsys4*), and the optimal value of (3.1.27, *BAP*) is given by

(*BAPN*)

$$\begin{aligned} \inf_{\lambda \in \Lambda} |\mu - \lambda|^2_{\mathcal{G}^*} &= |\mu - \sum_{k=1}^{M} \alpha_k(\mu)\lambda_k|^2_{\mathcal{G}^*} \\ &= (\mu, \mu)_{\mathcal{G}^*} - 2\sum_{k=1}^{M} \alpha_k(\mu)(\lambda_k, \mu)_{\mathcal{G}^*} \\ &\quad + \sum_{j,k=1}^{M} \alpha_j(\mu)\alpha_k(\mu)(\lambda_j, \lambda_k)_{\mathcal{G}^*}. \end{aligned} \tag{3.1.29}$$

3.1.6 Properties of Optimal Recoveries

Assume that we used the method of section 3.1.4 (*Necessity)* to recover an element $g \in \mathcal{G}$ by some element g^* that satisfies

(*ljg*)

$$\lambda_j(g) = \lambda_j(g^*),\ 1 \leq j \leq M \tag{3.1.30}$$

for a set of linearly independent functionals $\lambda_1, \ldots, \lambda_M$ with representers g_j in the sense of

$$\lambda_j(v) = \lambda_j(\Pi_{\mathcal{P}}(v)) + (v, g_j)_{\mathcal{G}},\ v \in \mathcal{G}.$$

Assume further that the sufficient condition for uniqueness holds, as given in Theorem 3.1.19 (*ORT2)*, and that we normalized the functions g_j to satisfy $g_j = R_{\mathcal{P}} g_j$ or $\Pi_{\mathcal{P}} g_j = 0$.

Since any element $g^* = p \in \mathcal{P}$ satisfies (3.1.15, *charmin*), we get

Theorem 3.1.31 *(PolRepT1) The optimal recovery process reproduces elements of* $\mathcal{P}$. □

Corollary 3.1.32 *(PolRepCol) If* g^* *is the unique optimal recovery of* g*, then* $\Pi_{\mathcal{P}}(g - g^*) = 0$.

Proof: If $p \in \mathcal{P}$ is arbitrary, then clearly $(g+p)^* = g^* + p$ due to uniqueness. The recovery process thus acts separately on the two parts of $\mathcal{G} = \mathcal{P} + R_{\mathcal{P}}(\mathcal{G})$ with values in the respective parts of $\mathcal{S} = \mathcal{P} + R_{\mathcal{P}}(\mathcal{S})$. But then $(R_{\mathcal{P}} g)^* = R_{\mathcal{P}}(g^*)$ holds and

$$R_{\mathcal{P}}(g^*) = (R_{\mathcal{P}} g)^* = (g - \Pi_{\mathcal{P}} g)^* = g^* - \Pi_{\mathcal{P}} g$$

implies $\Pi_{\mathcal{P}} g^* = \Pi_{\mathcal{P}} g$. □

Turning to orthogonality relations, we have

$$(g_j, g - g^*)_{\mathcal{G}} + \lambda_j \Pi_{\mathcal{P}}(g - g^*) = 0,\ 1 \leq j \leq M$$

and for each element s from the space

(*DefS*)

$$\mathcal{S} = \left\{ \sum_{j=1}^{M} \alpha_j g_j + \sum_{k=1}^{q} \beta_k p_k \ :\ \sum_{j=1}^{M} \alpha_j \lambda_j(\mathcal{P}) = \{0\} \right\} \tag{3.1.33}$$

we get the orthogonality

(*EqOrtho*)

$$(s, g - g^*)_{\mathcal{G}} = 0 \tag{3.1.34}$$

by summation. But this means that g^* is a best approximation to g from $\mathcal{S}$:

Theorem 3.1.35 *(ORTBA) The solution g^* of the optimal recovery problem (3.1.4, ORP) for data from some element $g \in \mathcal{G}$ is a best approximation to g from the space $\mathcal{S}$ of (3.1.33, DefS) in the sense*

$$|g - g^*|_{\mathcal{G}} = \min_{s \in \mathcal{S}} |g - s|_{\mathcal{G}}.$$

□

We now proceed towards the prototype of an error bound. We use the space (3.1.28, *DefL*) of functionals and (3.1.30, *ljg*) to get $\lambda(g - g^*) = 0$ for all $\lambda \in \Lambda$. Now take any $\mu \in \mathcal{G}^*$ and form

$$\mu(g-g^*) = (\mu-\lambda)(g-g^*) \leq |(\mu-\lambda)\Pi_{\mathcal{P}}(g-g^*)| + |(g_\mu - g_\lambda, g-g^*)_{\mathcal{G}}| \leq |\mu-\lambda|_{\mathcal{G}^*}|g-g^*|_{\mathcal{G}},$$

using Corollary 3.1.32 (*PolRepCol*).

Theorem 3.1.36 *(ORTFA) The solution g^* of the optimal recovery problem (3.1.4, ORP) for data from some element $g \in \mathcal{G}$ satisfies the error bound*

(Eq2inf)

$$|\mu(g - g^*)| \leq \inf_{\lambda \in \Lambda} |\mu - \lambda|_{\mathcal{G}^*} \inf_{s \in \mathcal{S}} |g - s|_{\mathcal{G}} \tag{3.1.37}$$

for any functional $\mu \in \mathcal{G}^*$. □

The crucial factor in the error bound (3.1.37, *Eq2inf*) is the **generalized optimal power function**

(*GPDef*)

$$P(\mu) := P_\Lambda(\mu) := \inf_{\lambda \in \Lambda} |\mu - \lambda|_{\mathcal{G}^*} \tag{3.1.38}$$

with Λ from (3.1.28, *DefL*). If the functionals λ_j are "near" to μ, this quantity should be rather small, and we shall prove specific bounds later in 3.5.8 (*hrhodef*). This is made possible by the representation for $P(\mu)$ that follows readily from (3.1.27, *BAP*) and (3.1.29, *BAPN*), and which will also be useful in section 3.4.6 (*URT*).

3.1.7 Remarks

The theory of optimal recovery starts with the early paper of Golomb and Weinberger [12](*golomb-weinberger:59-1*), while reproducing kernel Hilbert spaces are much older (see e.g. the textbook by Meschkowski [4](*Meschkowski*)). A milestone was the theory of optimal recovery in the sense of Micchelli, Rivlin, and Winograd ([20](*micchelli-rivlin:77-1*) [21](*micchelli-rivlin:78-1*) [22](*micchelli-rivlin:84-1*) [23](*micchelli-et-al:76-1*)), while the current extension into the direction of information-based complexity is in [4](*bojanov-wozniakowski:92-1*).

3.2 Spaces of Functions

(*SecSF*) In order to arrive at conditionally positive **functions**, we now have to specialize our results on optimal recovery to the case of optimal recovery of functions.

3.2.1 From Hilbert Spaces to Positive Definite Functions

(*SecHSPDF*) We now specialize to a Hilbert space $\mathcal{F}$ of functions on some domain Ω that we do not restrict. But since classical functions are objects that allow the action of specific linear functionals

(*deltadef*)

$$\delta_x \ : \ g \mapsto g(x), \ g \in \mathcal{F}, \ x \in \Omega \tag{3.2.1}$$

called point-evaluation functionals, we assume that the above functionals δ_x are in $\mathcal{F}^*$ and thus continuous on $\mathcal{F}$. Then one can invoke the Riesz representation theorem 10.3.14 (*RieszT*) to get a function $g_{\delta_x} \in \mathcal{F}$ for each $x \in \Omega$ such that

(*DRKF*)

$$g(x) = \delta_x(g) = (g, g_{\delta_x})_{\mathcal{F}} \tag{3.2.2}$$

holds for all $g \in \mathcal{F}$, $x \in \Omega$. We now define a function

$$\Phi \ : \ \Omega \times \Omega \to I\!R, \ \Phi(x,y) := (g_{\delta_x}, g_{\delta_y}), \ x, y \in \Omega$$

and get

Theorem 3.2.3 *(PDFT1) If the point evaluation functionals in a Hilbert space $\mathcal{F}$ of functions on some domain Ω are continuous, then the space has a* reproducing kernel function Φ *with the following properties:*

1. $\Phi \ : \ \Omega \times \Omega \to I\!R$,
2. $\Phi(x, \cdot) = \Phi(\cdot, x) \in \mathcal{F}$ *for all* $x \in \Omega$,
3. $g(x) = (g, \Phi(x, \cdot))_{\mathcal{F}}$ *for all* $g \in \mathcal{F}, \ x \in \Omega$.

Proof: By definition and (3.2.2, *DRKF*),

$$\begin{array}{rcccl} g_{\delta_y}(x) & = & (g_{\delta_y}, g_{\delta_x})_{\mathcal{F}} & = & \Phi(y,x) \\ (g_{\delta_y}, g_{\delta_x})_{\mathcal{F}} & = & (g_{\delta_x}, g_{\delta_y})_{\mathcal{F}} & = & \Phi(x,y) = \Phi(y,x) \end{array}$$

for all $x, y \in \Omega$, proving all of the assertions. □

We now compare this with Definition 1.5.3 (*DPD*) from section 1.5 (*subsecEIA*) on page 10 which we restate here for convenience:

Definition 3.2.4 *A real-valued function*

$$\Phi : \Omega \times \Omega \to I\!R$$

is a **positive definite function** *on Ω, iff for any choice of finite subsets $X = \{x_1, \ldots, x_M\} \subseteq \Omega$ of M different points the matrix*

$$A_{X,\Phi} = (\Phi(x_k, x_j))_{1 \le j,k \le M}$$

is positive definite.

To test the function Φ from Theorem 3.2.3 (*PDFT1)* for positive definiteness, consider a finite subset $X = \{x_1, \ldots, x_M\} \subseteq \Omega$ of M different points and take an arbitrary vector $\alpha \in I\!R^M$. Then

$$\alpha^T A_{X,\Phi} \alpha = \sum_{j,k=1}^{M} \alpha_j \alpha_k \Phi(x_k, x_j) = \left(\sum_{j=1}^{M} \alpha_j g_{x_j}, \sum_{k=1}^{M} \alpha_k g_{x_k} \right)_{\mathcal{F}} = \| \sum_{j=1}^{M} \alpha_j g_{x_j} \|_{\mathcal{F}}^2$$

implies that the matrix $A_{X,\Phi}$ always is positive semidefinite, because it is the Gramian of the functions $g_{x_j},\ 1 \le j \le M$. It is positive definite if and only if these functions are linearly independent in $\mathcal{F}$. Furthermore, is is easy to see from (3.2.2, *DRKF*) that the functions $g_{x_j},\ 1 \le j \le M$ are linearly dependent if and only if the point evaluation functionals $\delta_{x_j},\ 1 \le j \le M$ are linearly dependent in the dual space $\mathcal{F}$. Another simple exercise is to show equivalence of the linear independence of $\delta_{x_j},\ 1 \le j \le M$ with each of the following notions:

Definition 3.2.5 *(DFSP) A space $\mathcal{F}$ of functions on some domain Ω has the* **finite separation property***, if for all finite subsets $X = \{x_1, \ldots, x_M\} \subseteq \Omega$ of M different points there are M functions $g_1, \ldots, g_M \in \mathcal{F}$ that separate the points in $X = \{x_1, \ldots, x_M\}$, i.e.*

$$g_j(x_k) = \delta_{jk},\ 1 \le j,k \le M.$$

Definition 3.2.6 *(DFIP) A space $\mathcal{F}$ of functions on some domain Ω has the* **finite interpolation property***, if for all finite subsets $X = \{x_1, \ldots, x_M\} \subseteq \Omega$ of M different points and all vectors $\alpha \in I\!R^M$ there is a function $g \in \mathcal{F}$, depending on $X = \{x_1, \ldots, x_M\}$ and α, such that*

$$g(x_k) = \alpha_k,\ 1 \le k \le M.$$

We combine this into a result that proves the setting in 1.5 (*subsecEIA)* to occur naturally in fairly general situations:

Theorem 3.2.7 *Let $\mathcal{F}$ be a space of real-valued functions on some domain Ω, and assume*

1. *$\mathcal{F}$ is a Hilbert space over $\mathbb{R}$,*
2. *the point evaluation functionals (3.2.1, deltadef) are continuous on $\mathcal{F}$,*
3. *$\mathcal{F}$ has the finite interpolation or the finite separation property.*

Then $\mathcal{F}$ is a reproducing kernel Hilbert space, and its kernel function $\Phi : \Omega \times \Omega$ is a positive definite function. □

3.2.2 Generalization towards Conditionally Positive Definite Functions

(*SecGCPDF*) We now return to the slightly more general setting of section 3.1.1 (*subsecORP*). The continuous linear functionals now have to vanish on the kernel $\mathcal{P}$ of the bilinear form $(\cdot,\cdot)_{\mathcal{G}}$, and this is not a usual property of point evaluation functionals. But we can resort to the functionals

(*deltadef2*)

$$\delta_{x,\mathcal{P}} := \delta_x - \delta_x(\Pi_{\mathcal{P}}) \tag{3.2.8}$$

that will vanish on $\mathcal{P}$ for all $x \in \Omega$. We thus should require the functionals $\delta_{x,\mathcal{P}}$ from (3.2.8, *deltadef2*) to be continuous with respect to the bilinear form $(\cdot,\cdot)_{\mathcal{G}}$. This is the same as to assume that the point evaluation functionals δ_x are in $\mathcal{G}^*$, and then we can use (3.1.10, *lrep*) to get the generalization

(*DRKF2*)

$$\delta_{x,\mathcal{P}}(g) = g(x) - (\Pi_{\mathcal{P}}(g))(x) = (g, g_{\delta_{x,\mathcal{P}}})_{\mathcal{G}} \tag{3.2.9}$$

of (3.2.2, *DRKF*) for all $g \in \mathcal{G}$, $x \in \Omega$. This is a special form of (3.1.10, *lrep*) on page 31 and yields the Taylor-type formula

(*Taylor*)

$$g(x) = (\Pi_{\mathcal{P}}(g))(x) + (g, g_{\delta_{x,\mathcal{P}}})_{\mathcal{G}} \tag{3.2.10}$$

for all $g \in \mathcal{G}$, $x \in \Omega$. We now define

(*DefPhiGen*)

$$\Phi \ : \ \Omega \times \Omega \to \mathbb{R},\ \Phi(x,y) := (g_{\delta_{x,\mathcal{P}}}, g_{\delta_{y,\mathcal{P}}})_{\mathcal{G}},\ x,y \in \Omega \tag{3.2.11}$$

and get

Theorem 3.2.12 *(CPDFT1) If the functionals (3.2.8, deltadef2) for a space $\mathcal{G}$ of functions on some domain Ω are continuous with respect to the bilinear form $(\cdot,\cdot)_{\mathcal{G}}$ with finite-dimensional kernel $\mathcal{P}$ and projector $\Pi_{\mathcal{P}} : \mathcal{G} \to \mathcal{P}$, then the space has a* reproducing kernel function Φ *with the following properties:*

1. $\Phi : \Omega \times \Omega \to I\!R$,
2. $\Phi(x,\cdot) = \Phi(\cdot,x) \in \mathcal{G}$ *for all* $x \in \Omega$,
3. $\Pi_{\mathcal{P}}\Phi(x,\cdot) = \Pi_{\mathcal{P}}\Phi(\cdot,x) = 0$ *for all* $x \in \Omega$,
4. $\Phi(x,y) = (\Phi(x,\cdot),\Phi(y,\cdot))_{\mathcal{G}}$ *for all* $x,y \in \Omega$
5. $g(x) = \Pi_{\mathcal{P}}(g)(x) + (g,\Phi(x,\cdot))_{\mathcal{G}}$ *for all* $g \in \mathcal{G},\ x \in \Omega$.

Proof: We proceed exactly as in Theorem 3.2.3 (*PDFT1)* and get (*PhiRep2*)

$$\Phi(x,y) = (g_{\delta_{x,\mathcal{P}}}, g_{\delta_{y,\mathcal{P}}})_{\mathcal{G}} = g_{\delta_{x,\mathcal{P}}}(y) - (\Pi_{\mathcal{P}} g_{\delta_{x,\mathcal{P}}})(y). \tag{3.2.13}$$

This proves properties 2 and 3, while 1 holds by definition. Putting the above identity into (3.2.10, *Taylor*) and (3.2.11, *DefPhiGen*) yields the fourth and fifth property. □

We shall see later that the well-known conditionally positive definite functions fail to satisfy some of these properties, but there is a fairly standard process that shows how to get the properties by slight modifications. We shall comment on this when we consider the construction of native Hilbert spaces from given conditionally positive definite functions in section 3.3 (*SecNS).*

The identity (3.1.10, *lrep*) on page 31 introduced a representing function $g_\lambda \in \mathcal{G}$ for each functional $\lambda \in \mathcal{G}^*$. This was used in (3.1.11, *lrepj*) to derive the system (3.1.14, *EQsys3*) for solving the recovery problem. To bring this into line with the system (1.7.2, *EQsys2*) on page 14, we use (3.2.13, *PhiRep2*) to form

$$\begin{aligned} \lambda^y \Phi(x,y) &= \lambda(g_{\delta_{x,\mathcal{P}}}) - \lambda \Pi_{\mathcal{P}} g_{\delta_{x,\mathcal{P}}} \\ &= (g_\lambda, g_{\delta_{x,\mathcal{P}}})_{\mathcal{G}} = g_\lambda(x) - (\Pi_{\mathcal{P}} g_\lambda)(x) \end{aligned}$$

and get

$$g_\lambda = \Pi_{\mathcal{P}} g_\lambda + \lambda^y \Phi(\cdot,y)$$

for all $\lambda \in \mathcal{G}^*$. Since g_λ is nonunique modulo functions from $\mathcal{P}$, we even can omit the first summand and use the above equation as a definition for g_λ.

With a second functional $\mu \in \mathcal{G}^*$ we can write

$$\begin{aligned} \mu^x \lambda^y \Phi(x,y) &= \mu g_\lambda - \mu \Pi_{\mathcal{P}} g_\lambda \\ &= \mu \Pi_{\mathcal{P}} g_\lambda + (\mu,\lambda)_{\mathcal{G}^*} - \mu \Pi_{\mathcal{P}} g_\lambda \\ &= (\mu,\lambda)_{\mathcal{G}^*}. \end{aligned}$$

This proves

(*gjkrep*)

$$(g_j, g_k)_{\mathcal{G}} = (\lambda_j, \lambda_k)_{\mathcal{G}^*} = \lambda_j^x \lambda_k^y \Phi(x,y) \tag{3.2.14}$$

for the elements of the matrix in (3.1.14, *EQsys3*).

We now want to move towards conditionally positive definite functions, but we still have to replace polynomials in Definition 1.6.2 (*DCPD)* on page 12:

Definition 3.2.15 *(DCPD2) A real-valued function*

$$\Phi : \Omega \times \Omega \to I\!R$$

is a **conditionally positive definite function** *with respect to a finite-dimensional space $\mathcal{P}$ of functions on Ω, iff for any choice of finite subsets $X = \{x_1, \ldots, x_M\} \subseteq \Omega$ of M different points the value*

$$\alpha^T A_{X,\Phi} \alpha := \sum_{j,k=1}^{M} \alpha_j \alpha_k \Phi(x_j, x_k)$$

of the quadratic form (1.6.1, QFdef) is positive, provided that the vector $\alpha = (\alpha_1, \ldots, \alpha_M) \in I\!R^M \setminus \{0\}$ has the additional property

(CPDef2)

$$\sum_{j=1}^{M} \alpha_j p(x_j) = 0 \tag{3.2.16}$$

for all $p \in \mathcal{P}$.

Theorem 3.2.17 *(CPDNeccT) Let $\mathcal{G}$ be a space of real-valued functions on some domain Ω, and assume*

1. *$\mathcal{G}$ has a real-valued symmetric bilinear form $(\cdot,\cdot)_{\mathcal{G}}$ with a finite dimensional kernel $\mathcal{P}$ and corresponding projector $\Pi_{\mathcal{P}}$,*
2. *the point evaluation functionals (3.2.8, deltadef2) are continuous with respect to the bilinear form,*

3. $\mathcal{G}$ has the finite interpolation or the finite separation property.

Then $\mathcal{G}$ has a reproducing kernel in the sense of Theorem 3.2.12 (CPDFT1), and its kernel function $\Phi : \Omega \times \Omega$ is a conditionally positive definite function with respect to $\mathcal{P}$.

Proof: Again, we consider a finite subset $X = \{x_1, \ldots, x_M\} \subseteq \Omega$ of M different points, but now we take a vector $\alpha \in \mathbb{R}^M$ with (3.2.16, *CPDef2*). Then we can repeat the steps of the proof of Theorem 3.2.3 (*PDFT1)* to see that the matrix $A_{X,\Phi}$ is positive semidefinite. To prove definiteness, we now assume that

(*inP*)

$$\sum_{j=1}^{M} \alpha_j g_{\delta_{x_j},\mathcal{P}} \in \mathcal{P} \tag{3.2.18}$$

holds and have to prove that α is zero. But (3.2.16, *CPDef2*) and (3.2.18, *inP*) imply via (3.2.9, *DRKF2*) that the point evaluation functionals $\delta_{x_j},\ 1 \leq j \leq M$ are linearly dependent. □

We see that conditionally positive definite functions arise necessarily whenever optimal recovery of functions from a space $\mathcal{G}$ with a bilinear form is attempted. The coefficient matrix of the major part of the linear system has elements of Gramian form $(g_j, g_k)_{\mathcal{G}}$, even if the recovery is carried out in more general (non-function-) spaces. This means that positive (semi-) definiteness is the natural condition to ask for, and there is no reason to replace it by nonsingularity.

3.2.3 Sobolev and Beppo-Levi Spaces

We now want to exhibit some special cases where we can start from a space $\mathcal{G}$ with bilinear form and arrive at a conditionally positive definite function. The most usual bilinear form defined on functions is the L_2 inner product

$$(f, g)_{L_2(\Omega)} := \int_{x \in \Omega} f(x) g(x) dx$$

However, point evaluation functionals are not continuous with respect to this inner product. This is easy to see when looking at the evaluation at zero of functions of the form $f_\alpha(x) := \exp(-\alpha \|x\|_2^2)$ for large positive α. The $L_2(\mathbb{R}^d)$ inner products tend to zero for $\alpha \to \infty$, while the value at zero is always one. Thus there is no positive constant C such that

$$|\lambda(f_\alpha)| \leq C \|f_\alpha\|_{L_2(\mathbb{R}^d)}$$

holds. As a warm-up for similar calculations occurring in later sections of the text, let us do the evaluation of the inner product. It suffices to take $\beta = 2\alpha$ and calculate the integral

$$\int_{x\in\Omega} \exp(-\beta\|x\|_2^2)dx = \text{vol}\,(S^{d-1})\int_0^\infty r^{d-1}\exp(-\beta r^2)dr$$

by going over to polar coordinates and integrating over the scaled unit sphere $S^{d-1} \subset I\!R^d$. Its surface area (or its $d-1$-dimensional volume) is $\text{vol}\,(S^{d-1}) = 2\pi^{(d-1)/2}/\Gamma((d-1)/2)$ due to (10.4.2, *VolS*). The rest follows from substitution and the definition (10.4.1, *GammaDef*) of the Gamma function:

$$\begin{aligned}\int_0^\infty r^{d-1}\exp(-\beta r^2)dr &= \tfrac{1}{2\beta}\int_0^\infty (\frac{t}{\beta})^{d/2-1}\exp(-t)dt \\ &= \tfrac{1}{2}\beta^{-d/2}\Gamma(d/2).\end{aligned}$$

If the reader has difficulties with this, it is time to work through part 10.4 (*SecSFT)* of the appendix.

To make point evaluation functionals continuous, we require a stronger bilinear form than just the L_2 inner product. And the above discussion shows that problems may get worse with increasing space dimension.

The usual trick is to introduce derivatives into the bilinear form. In particular, take a multiindex $\alpha \in Z\!\!Z^d_{\geq 0}$ and define f^α as the multivariate derivative of order α of some function f. For a fixed integer $m \geq 0$, assemble all derivatives with $|\alpha| := \|\alpha\|_1 = m$ into a positive semidefinite bilinear form

$$(f,g)_m := \int_\Omega \sum_{|\alpha|=m} \left(\begin{array}{c} m \\ \alpha \end{array}\right) f^\alpha(x) g^\alpha(x) dx$$

on all functions that are at least in $C^m(\Omega)$. Here, we used the multivariate version of

$$\left(\begin{array}{c} m \\ \alpha \end{array}\right) := \frac{m!}{\alpha_1!\ldots,\alpha_d!} \text{ with } |\alpha| = m.$$

For simply connected domains Ω with a nonzero interior in $I\!R^d$ the nullspace of the bilinear form will then coincide with the space $\mathcal{P} = I\!P^d_m$ of polynomials of order m on $I\!R^d$. To do this, we need that a C^m function on Ω with vanishing derivatives of order m must necessarily be a polynomial, and this works nicely in the interior of Ω by application of the multivariate Taylor formula. The boundary does not count for the integral, and the polynomial is unique, if we do not have multiple components of the domain.

However, we still have to check the continuity of point-evaluation functionals $\delta_{x,\mathcal{P}}$ in the sense of (3.2.8, *deltadef2*) on page 42. The construction of a suitable projector $\Pi_{\mathcal{P}}$ to the nullspace $\mathcal{P} = I\!P_m^D$ will be given in Lemma 3.5.29 (*LemPIG)* on page 86 for use in a different context, but it is actually no big deal. Much more serious is the proof of the fact that $m > d/2$ is necessary and sufficient for continuity of the point-evaluation functionals. This is called the **Sobolev inequality**, but its proof is delayed to 10.5.3 (*SecSob)*.

If we assume $m > d/2$ and start with of $\mathcal{G} = C^m(\Omega)$ in the sense of section 3.2.2 (*SecGCPDF)*, we still have to form the Hilbert space completion and to derive the functions $g_{\delta_{x,\mathcal{P}}}$ that occur in (3.2.9, *DRKF2*) and allow to define a normalized conditionally positive definite function Φ via (3.2.11, *DefPhiGen*). To do these things on the full space $I\!R^d$ will later turn out to be much easier than to use a compact domain Ω. To avoid problems with nonexistence of $\|f\|_m$, we restrict ourselves to the subspace of $C^m(I\!R^d)$ of functions with bounded seminorm $|\cdot|_m$. The resulting completed space $\mathcal{G}$ with the bilinear form $(\cdot,\cdot)_m$ is called the **Beppo-Levi space** of order m on $I\!R^d$. For readers without a background in partial differential equations it will probably be a surprise to hear that the resulting Φ then precisely is the normalization of the conditionally positive definite radial function $\phi(r) = r^{2m-d}$ for d odd and $\phi(r) = r^{2m-d}\log r$ for d even.

We give a brief and sloppy "physicist-style" explanation for this and do the strict proof the other way round: we later construct the space from the conditionally positive definite function along the lines of the next section. The informal technique just takes (3.2.9, *DRKF2*) for granted and rewrites it in the form

$$\begin{aligned}\delta_{x,\mathcal{P}}(g) &= \int_\Omega \sum_{|\alpha|=m} \left(\begin{array}{c} m \\ \alpha \end{array}\right) g^\alpha(y) g^\alpha_{\delta_{x,\mathcal{P}}}(y)dy \\ &= (-1)^m \int_\Omega g(y) \sum_{|\alpha|=m} \left(\begin{array}{c} m \\ \alpha \end{array}\right) g^{2\alpha}_{\delta_{x,\mathcal{P}}}(y)dy\end{aligned}$$

if boundary terms are neglected. Thus, in the sense of linear partial differential equations, the function $g_{\delta_{x,\mathcal{P}}}$ must (up to a sign) be a fundamental solution corresponding to the differential operator

$$g \mapsto (-1)^m \sum_{|\alpha|=m} \left(\begin{array}{c} m \\ \alpha \end{array}\right) g^{2\alpha}$$

which (by a simple inductive proof) coincides with the m-th power $(-1)^m\Delta^m$ of the negative Laplacian

$$\Delta(f) := \sum_{j=1}^{d} \frac{\partial^2 f}{\partial x_j^2}.$$

This is the hidden reason for the $\begin{pmatrix} m \\ \alpha \end{pmatrix}$ factors in the definition of the bilinear form. The corresponding fundamental solutions are well-known and must be radial due to the radial symmetry of the Laplacian. Using the radial form of the Laplacian, they can be calculated explicitly, and they always are either of the form r^γ or $r^\gamma \log r$. The boundary conditions, when evaluated properly, force to take the solution with maximal smoothness in zero or with minimal decay at infinity, and this is the radial function given above.

The case $d = 2$ requires $m > d/2 = 1$, and the minimal possible m leads to $m = 2$ and $\phi(r) = r^2 \log r$. The corresponding differential operator is Δ^2, describing the surfaces formed by thin plates under external forces or constraints. This is where **thin-plate splines** have their name, and the original approach by Duchon started from the partial differential equation background of these functions. The other cases are fundamental solutions of the iterated Laplacian, and since solutions of the plain Laplacian are called **harmonic functions**, the radial functions of the form $\phi(r) = r^\beta$ for $\beta \notin 2\mathbb{Z}$ or $\phi(r) = r^\beta \log r$ for $\beta \in 2\mathbb{Z}$ are called **polyharmonic functions**. The transition to non-integer values of β is possible via Fourier transforms and will be done in general later.

Looking back at the seminorm $|\cdot|_m$ induced by the bilinear form $(\cdot,\cdot)_m$, we see that the optimal recovery problem attempts to pick a function with least weighted mean square of all derivatives of order m. This is somewhat like an energy minimization in case $m = 2$, but $m = 2$ is admissible only in spaces of dimension up to $d = 3$.

Another even more important space arises when all derivatives up to order m are summed up to generate a new bilinear form

$$((f,g))_m := \sum_{j=0}^{m} \sum_{|\alpha|=j} \int_\Omega \frac{\partial^\alpha f}{\partial x^\alpha} \frac{\partial^\alpha g}{\partial x^\alpha} dx$$

This is positive definite and defines via completion a Hilbert space $W_2^m(\Omega)$ called **Sobolev space** of order m. Again, the point evaluation functionals are continuous only if $m > d/2$ holds. Using Fourier transforms, the special

case $\Omega = I\!R^d$ can be treated explicitly and yields a positive definite radial basis function

$$\phi(r) = r^{m-d/2} K_{m-d/2}(r)$$

up to a factor depending on m and d, where K_ν is the Bessel or Macdonald function defined in (10.4.3, *KnuDef*). The power of r cancels the singularity of $K_{m-d/2}$ at zero exactly, since the asymptotics near zero are given by (10.4.4, *KnuAsyZero*).

These radial basis functions look strange, but they arise very naturally, Since the Bessel functions K_ν have exponential decay towards infinity due to (10.4.5, *KnuAsyInf*), the translates of $\phi(\|x\|_2)$ lead to virtually band-limited interpolation matrices. The evaluation of such functions is easily possible by calling standard subroutine packages.

If one considers other (equivalent) inner products on Sobolev spaces, the associated positive definite functions Φ will change. Naively, we would not expect these changes to be substantial, but surprisingly there is an equivalent inner product that generates a compactly supported radial basis function. We shall see this when we check the functions introduced by Wendland in [30](*wendland:95-1*).

3.2.4 Invariance Principles

(*SecIP*) The preceding discussion showed that conditionally positive definite functions associated to function spaces on $I\!R^d$ often come out to be radial. We shall now look at this phenomenon in more detail.

Assume that the domain Ω allows a group $\mathcal{T}$ of geometric transformations, and that the bilinear form $(\cdot,\cdot)_{\mathcal{G}}$ of the space $\mathcal{G}$ is invariant under transformations from $\mathcal{T}$. By this we mean the properties

(*GInv*)

$$\begin{array}{rcl} g \circ T & \in & \mathcal{G} \\ (f \circ T, g \circ T)_{\mathcal{G}} & = & (f,g)_{\mathcal{G}} \\ (\Pi_{\mathcal{P}} g) \circ T & = & \Pi_{\mathcal{P}}(g \circ T) \end{array} \tag{3.2.19}$$

for all $T \in \mathcal{T}$ and all $f, g \in \mathcal{G}$. Then there are two ways to interpret the action of a functional δ_{Tx} for $x \in \Omega$ and $T \in \mathcal{T}$:

$$\begin{array}{rcllll} \delta_{Tx}(g) & = & g(Tx) & = & (\Pi_{\mathcal{P}} g)(Tx) + (g, g_{\delta_{Tx}})_{\mathcal{G}} \\ & = & (g \circ T)(x) & = & (\Pi_{\mathcal{P}}(g \circ T))(x) + (g \circ T, g_{\delta_x})_{\mathcal{G}} \\ & & & = & (\Pi_{\mathcal{P}} g)(Tx) + (g \circ T, g_{\delta_x} \circ T^{-1} \circ T)_{\mathcal{G}} \\ & & & = & (\Pi_{\mathcal{P}} g)(Tx) + (g, g_{\delta_x} \circ T^{-1})_{\mathcal{G}} \end{array}$$

and this proves

$$g_{\delta_{Tx}} - g_{\delta_x} \circ T^{-1} \in \mathcal{P}$$

for all $g \in \mathcal{G},\ T \in \mathcal{T}$. But this can be inserted into the definition of Φ to get

$$\Phi(Tx, Ty) = (g_{\delta_{Tx}}, g_{\delta_{Ty}})_{\mathcal{G}} = (g_{\delta_x} \circ T^{-1}, g_{\delta_y} \circ T^{-1})_{\mathcal{G}} = (g_{\delta_x}, g_{\delta_y})_{\mathcal{G}} = \Phi(x, y)$$

for all $x, y \in \Omega$. We thus have

Theorem 3.2.20 *(InvT1) Let $\mathcal{G}$ and Φ satisfy the assumptions of Theorem 3.2.12 (CPDFT1). If the domain Ω allows a group $\mathcal{T}$ of transformations that leave the bilinear form $(\cdot,\cdot)_{\mathcal{G}}$ on $\mathcal{G}$ invariant in the sense of (3.2.19, GInv), then Φ is invariant under $\mathcal{T}$ in the sense*

(PhiInv)

$$\Phi(x, y) = \Phi(Tx, Ty) \tag{3.2.21}$$

for all $x, y \in \Omega,\ T \in \mathcal{T}$. □

Corollary 3.2.22 *If the domain Ω has a fixed element denoted by x_0, and if for all $x \in \Omega$ there is a transformation $T_x \in \mathcal{T}$ with $T_x(x) = x_0$, then Φ takes the form*

(Phi1arg)

$$\Phi(x, y) = \Phi(T_y(x), x_0) \tag{3.2.23}$$

such that one of the two arguments of Φ is redundant.

We now consider some examples of domains with groups of transformations, and we always assume the invariance requirements of Theorem 3.2.20 (*InvT1)* to be satisfied.

Example 3.2.24 *If Ω is itself a group with neutral element 1, then*

$$\Phi(x, y) = \Phi(y^{-1}x, 1)$$

for all $x, y \in \Omega$.

Example 3.2.25 *If $\Omega = I\!R^d$ with the group of translations, then*

(PhiDiff)

$$\Phi(x, y) = \Phi(y - x, 0) = \Phi(x - y, 0) \tag{3.2.26}$$

for all $x, y \in I\!R^d$.

Example 3.2.27 *If $\Omega = \mathbb{R}^d$ with the group of Euclidean rigid-body transformations (i.e. translations and rotations), then Φ is a radial function*

$$\Phi(x, y) = \phi(\|y - x\|_2)$$

for all $x, y \in \mathbb{R}^d$, where $\phi : \mathbb{R}_{\geq 0} \to \mathbb{R}$.

Proof: First use the translations of the previous case to write $\Phi(x, y) = \Phi(x - y, 0)$, and then rotate $x - y$ to a fixed unit vector in $\mathbb{R}^d$ multiplied by $\|x - y\|_2$. Then we are left with a scalar function of $\|x - y\|_2$. □

We note the remarkable fact that conditionally positive definite radial basis functions always occur in optimal recovery problems on $\mathbb{R}^d$ for functions from spaces that carry a bilinear form with Euclidean invariance.

Example 3.2.28 *If $\Omega = S^{d-1} \subset \mathbb{R}^d$ is the $(d-1)$-sphere, i.e. the surface of the unit ball in $\mathbb{R}^d$, then rotational invariance implies that Φ is* **zonal**, *i.e.*

$$\Phi(x, y) = \phi(x^T y)$$

for all $x, y \in S^{d-1}$, where $\phi : [0, 1] \to \mathbb{R}$.

In this case the function Φ can be written as a scalar function of the angle between the two arguments, or the cosine of this angle.

Example 3.2.29 *If $\Omega = \mathbb{R}^d$ and if the group $\mathcal{T}$ is $\mathbb{Z}^d$ under addition, then $\mathcal{G}$ is a shift-invariant space (see [5](boor-et-al:94-2)), and Φ is fully determined by its values on $\mathbb{R}^d_{\geq 0} \times \mathbb{R}^d_{\geq 0}$.*

In this case, pick T to shift $\lfloor \min(x, y) \rfloor$ to the origin, using minimum and $\lfloor \cdot \rfloor$ coordinatewise.

Example 3.2.30 *If $\Omega = [-\pi, \pi]^d$, if the space $\mathcal{G}$ consists of d-variate 2π-periodic functions, and if the bilinear form is invariant under coordinatewise real-valued shifts, then we are in a fully periodic setting and $\Phi(x, y)$ has the form (3.2.26, PhiDiff) with a 2π-periodic first argument.*

3.2.5 Remarks

The monograph [2](*atteia:92-1*) also explores the relation between reproducing kernel Hilbert spaces and associated recovery problems. This section used parts of [8](*RSTranslInv*).

3.3 Native Spaces

(*SecNS*) The previous sections have shown that each Hilbert space setting of a recovery problem leads to a specific conditionally positive definite function acting as a reproducing kernel. We now turn this upside down: for each conditionally positive definite function Φ there is a Hilbert space with reproducing kernel Φ. This seems to be a quite academic question, but it isn't. The main reason is that it is much more easy to construct useful conditionally positive definite functions than to find certain Hilbert spaces. Thus it often happens that one starts with a conditionally positive definite function, not with a Hilbert space. Furthermore, if a conditionally positive definite function Φ is constructed without any relation to a Hilbert space, the latter can be theoretically defined and nicely used to investigate the recovery quality of Φ..

3.3.1 From Conditionally Positive Definite Functions to Hilbert Spaces

Now let Φ be a conditionally positive definite function on some domain Ω with respect to some finite-dimensional space $\mathcal{P}$ in the sense of Definition 3.2.15 (*DCPD2)* on page 44. We have to construct the space $\mathcal{G}$ occurring the preceding sections, and its associated bilinear form with nullspace $\mathcal{P}$. Since there is no other tool available than the definition of conditionally positive definite functions, we first have to work with finitely supported functionals

(*Deflxma*)

$$\lambda_{X,M,\alpha} \ : \ f \mapsto \sum_{j=1}^{M} \alpha_j f(x_j) \tag{3.3.1}$$

for arbitrary subsets $X = \{x_1, \ldots, x_M\} \subset \Omega$ of M distinct points, where the coefficient vector $\alpha \in \mathbb{R}^M$ satisfies (3.2.16, *CPDef2*), i.e. the above functional is zero on the space $\mathcal{P}$. We thus define $\mathcal{P}_\Omega^\perp$ to be the set containing all of these functionals. To turn $\mathcal{P}_\Omega^\perp$ into a vector space over $\mathbb{R}$, we use the obvious multiplication by scalars and define the sum of $\lambda_{X,M,\alpha}$ and $\lambda_{Y,N,\beta}$ as $\lambda_{Z,L,\gamma}$ with $Z = \{z_1, \ldots, z_L\}$ and

$$\begin{array}{rcll} Z & = & X \cup Y & \\ L & = & \text{card}\,(Z) & \\ \gamma_\ell & = & \alpha_j & \text{if} \quad z_\ell = x_j \in X \setminus (X \cap Y) \\ \gamma_\ell & = & \beta_k & \text{if} \quad z_\ell = y_k \in Y \setminus (X \cap Y) \\ \gamma_\ell & = & \alpha_j + \beta_k & \text{if} \quad z_\ell = x_j = y_k \in X \cap Y. \end{array}$$

This definition makes sure that

$$\lambda_{X,M,\alpha}(f) + \lambda_{Y,N,\beta}(f) = \lambda_{Z,L,\gamma}(f)$$

holds for each function f on Ω, and thus the sum satisfies (3.2.16, *CPDef2*). The usual laws for vector spaces are satisfied, and we now define a bilinear form on $\mathcal{P}_\Omega^\perp$ by

(*DefBil*)

$$(\lambda_{X,M,\alpha}, \lambda_{Y,N,\beta})_\Phi := \sum_{j=1}^{M}\sum_{k=1}^{N} \alpha_j \beta_k \Phi(x_j, y_k). \tag{3.3.2}$$

Since Φ is positive definite with respect to $\mathcal{P}$, we even have positive definiteness of the bilinear form on $\mathcal{P}_\Omega^\perp$, and $\mathcal{P}_\Omega^\perp$ is a pre-Hilbert space with the inner product $(\cdot,\cdot)_\Phi$ introduced by Φ. Note that the vector space $\mathcal{P}_\Omega^\perp$ is only dependent on Ω and $\mathcal{P}$, not on Φ itself, but the inner product on $\mathcal{P}_\Omega^\perp$ depends on Φ, as we indicate by our notation.

We now can define the **native space** $\mathcal{G}$ with respect to Φ to consist of all functions on Ω on which all functionals from $\mathcal{P}_\Omega^\perp$ are continuous:

(*calgdef*)

$$\mathcal{G} := \left\{ f \ : \ \Omega \to I\!R, \ |\lambda(f)| \le C_f \|\lambda\|_\Phi \text{ for all } \lambda \in \mathcal{P}_\Omega^\perp \right\}. \tag{3.3.3}$$

It is immediately clear that $\mathcal{P}$ is a subset of $\mathcal{G}$, but it is neither clear nor true (in general) that the functions $\Phi(x,\cdot)$ are in $\mathcal{G}$. Furthermore, we still need a bilinear form on $\mathcal{G}$ that has $\mathcal{P}$ as its nullspace. To do this, we first define the map

$$F \ : \ \mathcal{P}_\Omega^\perp \to \mathcal{G}, \ F(\lambda_{X,M,\alpha}) = \sum_{j=1}^{M} \alpha_j \Phi(x_j, \cdot)$$

and have to make sure that the image is indeed in $\mathcal{G}$. But this follows from the very important identity

(*lmF*)

$$\lambda_{Y,N,\beta}(F(\lambda_{X,M,\alpha})) = (\lambda_{X,M,\alpha}, \lambda_{Y,N,\beta})_\Phi = \lambda_{X,M,\alpha}(F(\lambda_{Y,N,\beta})) \tag{3.3.4}$$

for all $\lambda_{X,M,\alpha}, \ \lambda_{Y,N,\beta} \in \mathcal{P}_\Omega^\perp$. Then we define $\mathcal{F}_0 := F(\mathcal{P}_\Omega^\perp)$ and assert

Lemma 3.3.5 *The sum $\mathcal{P} + \mathcal{F}_0$ is direct, and the map F is bijective.*

Proof: Indeed if $F(\lambda) = p \in \mathcal{P}$, then for all $\mu \in \mathcal{P}_\Omega^\perp$ we have $\mu(F(\lambda)) = \mu(p) = (\lambda, \mu)_\Phi = 0$ due to (3.3.4, *lmF*), proving both assertions at the same time. □

In the above proof we used shorthand notation for functionals in $\mathcal{P}_\Omega^\perp$, and we shall only return to the full notation if absolutely necessary.

We now can define an inner product on $\mathcal{F}_0$ via F, turning F into an isometry and $\mathcal{F}_0$ into a pre-Hilbert space:

$$(F(\lambda), F(\mu))_\Phi := (\lambda, \mu)_\Phi$$

for all $\lambda, \mu \in \mathcal{P}_\Omega^\perp$. We used the same notation for the inner product, since there will be no confusion between spaces of functions and functionals, respectively.

The next step is to go over to Hilbert space completions of $\mathcal{P}_\Omega^\perp$ and $\mathcal{F}_0$ in the sense of Theorem 10.3.11 (*HSCT)*. Then we get a continuous extension of the isometry F to the completions for free, and we denote this map again by F. The completion of $\mathcal{F}_0$ will be denoted by $\mathcal{F}$, and our final goal is to prove the validity of a direct sum like

(*GPF2*)

$$\mathcal{G} = \mathcal{P} + \mathcal{F} \tag{3.3.6}$$

to recover (3.1.8, *GPF1*) on page 30. But this is a hard task since we do not know that the elements of the completion $\mathcal{F}$ of $F(\mathcal{P}_\Omega^\perp)$ are functions on Ω at all, let alone that they lie in $\mathcal{G}$. However, we know that an abstract element f of $\mathcal{F}$ allows the action of all functionals $\lambda_{X,M,\alpha} \in \mathcal{P}_\Omega^\perp$, since (3.3.4, *lmF*) yields

(*lfgeneral*)

$$\lambda_{X,M,\alpha}(f) = (\lambda_{X,M,\alpha}, F^{-1}(f))_\Phi. \tag{3.3.7}$$

This immediately implies a proper definition of function values for f in case of $\mathcal{P} = \{0\}$, since we can define

(*lfsimple*)

$$f(x) := \lambda_{\{x\},1,1}(f) \tag{3.3.8}$$

for all $x \in \Omega$. This definition is consistent with what we know for functions in $\mathcal{F}_0$, and we could proceed to prove (3.3.6, *GPF2*). But we need a little detour for the case $\mathcal{P} \neq \{0\}$, since the above point evaluation functionals are not in $\mathcal{P}_\Omega^\perp$. To facilitate this, we again require a projector $\Pi_\mathcal{P}$ onto $\mathcal{P}$ as in section 3.1.2 (*SecHSP)* on page 29. We could copy this definition, but since we are in a space of functions now, we want to give a specific construction that can be expressed in terms of function values.

To get such a special projector, we shall assume the existence of a subset

$$\Xi = \{\xi_1, \ldots, \xi_r\} \subseteq \Omega$$

which is nondegenerate with respect to $\mathcal{P}$ and assume without loss of generality that Ξ has a minimal number r of distinct points. Then there is a standard argument from linear algebra that allows to conclude that r equals the dimension of $\mathcal{P}$. In fact, the map

$$p \mapsto (p(\xi_1), \ldots, p(\xi_r))^T \in I\!R^r$$

is injective and we have $q := \dim\ \mathcal{P} \leq r$. If $p_1, \ldots, p_q$ form a basis of $\mathcal{P}$, we can write down the injective $r \times q$ matrix

(*Prq*)

$$P := (p_k(\xi_j))_{1 \leq j \leq r,\, 1 \leq k \leq q} \tag{3.3.9}$$

and pick a subset of rows that generate a submatrix of maximal row rank. If this were a proper subset, we could reduce r by going over to a subset of Ξ. Thus P has maximal row rank r. But then we must have $q = r$, because there cannot be r linearly independent vectors in a space of dimension $q < r$.

This shows that we can assume $r = q = \dim \mathcal{P}$ and nonsingularity of the $q \times q$ matrix P of (3.3.9, *Prq*). We use this to go over to a Lagrange-type basis of $\mathcal{P}$ with respect to Ξ which we again denote by $p_1, \ldots, p_q$. Then P is the identity matrix and we can write every function $p \in \mathcal{P}$ as

(*PRq2*)

$$p(\cdot) = \sum_{j=1}^{q} p(\xi_j) p_j(\cdot). \tag{3.3.10}$$

This now yields the explicit form of a projector $\Pi_{\mathcal{P}}$ onto $\mathcal{P}$ as

$$\Pi_{\mathcal{P}}(f)(\cdot) := \sum_{j=1}^{q} f(\xi_j) p_j(\cdot)$$

for all functions that are at least defined on Ξ. The projector has the additional property

$$(f - \Pi_{\mathcal{P}} f)(\Xi) = \{0\}$$

for all functions f that are defined on Ξ, because of $\delta_{\xi_j, \Xi} = 0,\ 1 \leq j \leq q$. Note that $\pi_j(f) = f(\xi_j)$ holds if we compare (3.3.10, *PRq2*) with (3.1.5, *DefPN*).

So the projector is well-defined, but we cannot use it right away, since we first need nice functionals in $\mathcal{P}_\Omega^\perp$. But such functionals come from the projector via

(*deltagen*)

$$\delta_{x,\Xi}(f) := f(x) - (\Pi_{\mathcal{P}}(f))(x) = f(x) - \sum_{j=1}^{q} f(\xi_j)p_j(x) \tag{3.3.11}$$

for all $x \in \Omega$ and they annihilate $\mathcal{P}$, as required.

The notation $\delta_{x,\mathcal{P}}$ from (3.2.8, *deltadef2*) is very similar, but there will be no possible confusion. Similar variations of point evaluation functionals will occur later. These functionals are useful to prove an intermediate result that will be of some use later:

Lemma 3.3.12 *(SuffPol) If the action of all functionals λ from $\mathcal{P}_\Omega^\perp$ is zero on a given function f from $\mathcal{G}$, then f coincides with a function from $\mathcal{P}$ on Ω.*

Proof. : Just take the functionals $\delta_{x,\Xi}$ for all $x \in \Omega$, and look at

$$0 = \delta_{x,\Xi}(f) = f(x) - (\Pi_{\mathcal{P}}(f))(x).$$

□

We now could generalize (3.3.8, *lfsimple*) using the above functionals in (3.3.7, *lfgeneral*):

(*lf2*)

$$f(x) := (\delta_{x,\Xi}, F^{-1}(f))_\Phi,\ x \in \Omega,\ f \in \mathcal{F}. \tag{3.3.13}$$

This assigns specific function values to the abstract element of the closure $\mathcal{F}$ of $\mathcal{F}_0$. The assignment has the consequence that $f(\Xi) = \{0\}$ due to $\delta_{\xi_j,\Xi} = 0,\ 1 \le j \le q$, and thus it is rather an assignment of values to $f - \Pi_{\mathcal{P}} f$ than to f itself. We thus avoid this complication and define a mapping

$$R_{\mathcal{P}} \ :\ \mathcal{F} \to \mathcal{G}$$

by

(*lf3*)

$$(R_{\mathcal{P}} f)(x) := (\delta_{x,\Xi}, F^{-1}(f))_\Phi,\ x \in \Omega,\ f \in \mathcal{F}. \tag{3.3.14}$$

We have to show that this maps into $\mathcal{G}$, and for this we have to evaluate

$$\begin{aligned}\lambda_{X,M,\alpha}(R_{\mathcal{P}}(f)) &= \sum_{j=1}^{M}\alpha_j(\delta_{x_j,\Xi},F^{-1}(f))_\Phi \\ &= \left(\sum_{j=1}^{M}\alpha_j\delta_{x_j,\Xi},F^{-1}(f)\right)_\Phi .\end{aligned}$$

Now the functional in the bilinear form boils down to

$$\begin{aligned}\sum_{j=1}^{M}\alpha_j\delta_{x_j,\Xi}(f) &= \sum_{j=1}^{M}\alpha_j\left(f(x_j)-\sum_{k=1}^{q}f(\xi_k)p_k(x_j)\right)\\ &= \sum_{j=1}^{M}\alpha_j f(x_j)-\sum_{k=1}^{q}f(\xi_k)\sum_{j=1}^{M}\alpha_j p_k(x_j)\\ &= \sum_{j=1}^{M}\alpha_j f(x_j)-0\\ &= \lambda_{X,M,\alpha}(f),\end{aligned}$$

and we end up with

(*RfDef*)

$$\lambda_{X,M,\alpha}(R_{\mathcal{P}}(f)) = (\lambda_{X,M,\alpha},F^{-1}(f))_\Phi \tag{3.3.15}$$

which proves $R_{\mathcal{P}}(f)\in\mathcal{G}$.

Theorem 3.3.16 *(GPFT2) The spaces $\mathcal{P}$, $\mathcal{G}$, and $\mathcal{F}$ of functions on Ω form a direct sum*

$$\mathcal{G}=\mathcal{P}+R_{\mathcal{P}}(\mathcal{F}),$$

and $R_{\mathcal{P}}$ defined by (3.3.14, lf3) is an isometry between $\mathcal{F}$ and $R_{\mathcal{P}}(\mathcal{F})\subseteq\mathcal{G}$. The inner products on $\mathcal{F}$ and $R_{\mathcal{P}}(\mathcal{F})$ introduce a bilinear form

$$(g,h)_{\mathcal{G}} := (R_{\mathcal{P}}^{-1}(g-\Pi_{\mathcal{P}}g),R_{\mathcal{P}}^{-1}(h-\Pi_{\mathcal{P}}h))_{\mathcal{F}}$$

with nullspace $\mathcal{P}$ on $\mathcal{G}$.

Proof: The intersection of $\mathcal{P}$ and $R_{\mathcal{P}}(\mathcal{F})$ is zero, because the second space consists of functions vanishing on Ξ, and the only such function in the first space is the zero function. Thus the sum is direct, and we have to show that the sum fills all of $\mathcal{G}$. Before we do that, we take a look at the mapping $R_{\mathcal{P}}$ and check the topology of $\mathcal{G}$. Each function f in $\mathcal{G}$ has the well-defined norm

$$\|f\|_{\mathcal{G}} := \sup_{\lambda\in\mathcal{P}_\Omega^\perp\setminus\{0\}}\frac{|\lambda(f)|}{\|\lambda\|_\Phi},$$

and the identity (3.3.15, *RfDef*) immediately yields

$$\|R_{\mathcal{P}}(f)\|_{\mathcal{G}} = \|F^{-1}(f)\|_{\Phi} = \|f\|_{\Phi}$$

for all $f \in \mathcal{F}$. Thus $R_{\mathcal{P}}$ is isometric, and $R_{\mathcal{P}}(\mathcal{F})$ is the closure of $R_{\mathcal{P}}(\mathcal{F}_0)$ in $\mathcal{G}$.

We now proceed to show that $\mathcal{P} + R_{\mathcal{P}}(\mathcal{F})$ fills all of $\mathcal{G}$, and we shall construct the inverse of $R_{\mathcal{P}}$. Take an arbitrary function $f \in \mathcal{G}$ and define a functional L_f on the space $\mathcal{P}_{\Omega}^{\perp}$ by

$$L_f(\lambda) := \lambda(f),\ \lambda \in \mathcal{P}_{\Omega}^{\perp}.$$

This functional is continuous on $\mathcal{P}_{\Omega}^{\perp}$ because f is in $\mathcal{G}$, and it has a continuous extension to the closure of $\mathcal{P}_{\Omega}^{\perp}$ which is a space isomorphic to the Hilbert space $\mathcal{F}$. We thus invoke the Riesz represenation theorem 10.3.14 (*RieszT)* to get an element $S(f) \in \mathcal{F}$ with

$$L_f(\lambda) = \lambda(f) = (\lambda, F^{-1}(S(f)))_{\Phi} = (F(\lambda), S(f))_{\Phi} \text{ for all } \lambda \in \mathcal{P}_{\Omega}^{\perp}.$$

Using (3.3.15, *RfDef*), this turns into

$$\lambda_{X,M,\alpha}(R_{\mathcal{P}}S(f)) = (\lambda_{X,M,\alpha}, F^{-1}Sf)_{\Phi} = \lambda_{X,M,\alpha}(f)$$

and Lemma 3.3.12 (*SuffPol)* implies that $f - R_{\mathcal{P}}S(f)$ coincides with a function from $\mathcal{P}$ on Ω, and since $\Pi_{\mathcal{P}}R_{\mathcal{P}}$ is the zero mapping, we see that

$$f = \Pi_{\mathcal{P}}F + R_{\mathcal{P}}Sf$$

holds for all f in $\mathcal{G}$, proving that the direct sum fills all of $\mathcal{G}$. The statement on the bilinear form is straightforward to prove. □

To write down a more explicit representation of the functions from $\mathcal{G}$, we apply F to $\delta_{x,\Xi}$ and get

$$F(\delta_{x,\Xi})(\cdot) = \Phi(x,\cdot) - \sum_{k=1}^{q} \Phi(\xi_k,\cdot)p_k(x) \in \mathcal{G}.$$

Then (3.3.14, *lf3*) and Theorem 3.3.16 (*GPFT2)* imply the representation (*Taylor2*)

$$f(x) \;=\; \sum_{j=1}^{q} f(\xi_j)p_j(x) + \left(\Phi(x,\cdot) - \sum_{j=1}^{q} \Phi(\xi_j,\cdot)p_j(x), f(\cdot)\right)_{\Phi}, \qquad (3.3.17)$$

but note that the sum in the first argument of the bilinear form cannot easily be taken out, because $\Phi(x,\cdot)$ may not be in $\mathcal{G}$. The same problem prevents us from concluding that Φ serves as a reproducing kernel in the strong sense of Theorem 3.2.12 (*CPDFT1)*. A good candidate, however, is the readily available function

(*EqPsiDef*)

$$\begin{aligned}\Psi(x,y) &:= (\delta_{x,\Xi},\delta_{y,\Xi})_\Phi \\ &= (F(\delta_{x,\Xi}),F(\delta_{y,\Xi}))_\Phi,\end{aligned} \tag{3.3.18}$$

because then (3.3.17, *Taylor2*) yields

$$\begin{aligned}F(\delta_{y,\Xi})(x) &= (\Pi_\mathcal{P}F(\delta_{y,\Xi}))(x) + (F(\delta_{x,\Xi}),F(\delta_{y,\Xi}))_\Phi\,, \\ &= (\Pi_\mathcal{P}F(\delta_{y,\Xi}))(x) + \Psi(x,y)\end{aligned}$$

such that

$$\Psi(x,y) = F(\delta_{y,\Xi})(x) - (\Pi_\mathcal{P}F(\delta_{y,\Xi}))(x)$$

holds, proving that $\Psi(\cdot,x)$ is indeed in $\mathcal{G}$ and satisfies $\Pi_\mathcal{P}(\Psi(\cdot,y)) = 0$ for all $y\in\Omega$. The above identity can now be put into (3.3.17, *Taylor2*) to get (3.2.10, *Taylor*) via

(*Taylor4*)

$$\begin{aligned}f(x) &= (\Pi_\mathcal{P}f)(x) + (\Psi(x,\cdot)+\Pi_\mathcal{P}F(\delta_{x,\Xi}))(\cdot),f(\cdot))_\Phi\,, \\ &= (\Pi_\mathcal{P}f)(x) + (\Psi(x,\cdot),f(\cdot))_\Phi\,.\end{aligned} \tag{3.3.19}$$

Thus the function Ψ satisfies all we need for Theorem 3.2.12 (*CPDFT1)*, but we still have to look at its relation to the original function Φ:

(*DefSymmPsi*)

$$\begin{aligned}\Psi(x,y) &= F(\delta_{x,\Xi})(y) - (\Pi_\mathcal{P}F(\delta_{x,\Xi}))(y) \\ &= \Phi(x,y) - \sum_{j=1}^q \Phi(\xi_j,y)p_j(x) - (\Pi_\mathcal{P}(\Phi(x,\cdot) - \sum_{j=1}^q \Phi(\xi_j,\cdot)p_j(x)))(y) \\ &= \Phi(x,y) - \sum_{j=1}^q \Phi(\xi_j,y)p_j(x) - \sum_{k=1}^q \Phi(x,\xi_k)p_k(y) + \sum_{j,k=1}^q \Phi(\xi_k,\xi_j)p_k(x)p_j(y).\end{aligned} \tag{3.3.20}$$

Inspection of this equation and comparison with (3.3.2, *DefBil*) implies that Φ and Ψ generate the same bilinear form for the definition of the native space. Thus Ψ is also conditionally positive definite and the native spaces generated by Φ and Ψ coincide.

We can now add up the results of this section:

Theorem 3.3.21 *(CPDSuffT) Let Φ be a conditionally positive definite function on some domain Ω with a finite-dimensional nullspace $\mathcal{P}$ of functions on Ω that allows an interpolatory projector*

$$(\Pi_{\mathcal{P}} f)(\cdot) = \sum_{j=1}^{q} f(\xi_j) p_j(\cdot)$$

where $p_1, \ldots, p_q$ are a basis of $\mathcal{P}$ and $\xi_1, \ldots, \xi_q$ form a $\mathcal{P}$-nondegenerate subset Ξ of Ω. Then there is a **native space** *$\mathcal{G}$ for Φ, carrying a bilinear form with nullspace $\mathcal{P}$ and having the function Ψ as defined in (3.3.20, DefSymmPsi) as a reproducing kernel in the sense of Theorem 3.2.17 (CPDNeccT). The native space is formed by adding a Hilbert space to $\mathcal{P}$.* □

The transition from a conditionally positive definite function Φ to the function Ψ with (3.3.20, *DefSymmPsi*) will be called **normalization** in the sequel. We note that the normalized function Ψ can also be defined if the projector is not interpolatory, but rather of the more general form (3.1.5, *DefPN*).

3.3.2 Normalization of conditionally positive definite functions

(*PhiNormalization*) With the notation of the preceding section it is fairly easy to describe the reduction of a conditionally positive definite function to an unconditionally positive definite function. This process coincides with the normalization by (3.3.20, *DefSymmPsi*).

Theorem 3.3.22 *(RedCPDFT) Let Φ be a conditionally positive definite function with respect to the nullspace $\mathcal{P}$ of the bilinear form on $\mathcal{G}$, and let the projector $\Pi_{\mathcal{P}}$ onto $\mathcal{P}$ be interpolatory with a minimal $\mathcal{P}$-nondegenerate set $\Xi = \{\xi_1, \ldots, \xi_q\}$ of points of Ω. Then the normalized function Ψ defined as in (3.3.20, DefSymmPsi) is unconditionally positive definite on $\Omega \setminus \Xi$.*

Proof: Consider a finite subset $X = \{x_1, \ldots, x_M\}$ of $\Omega \setminus \Xi$ and an arbitrary coefficient vector $\alpha \in \mathbb{R}^M$. Then the functional

$$\begin{aligned}\left(\sum_{j=1}^{M} \alpha_j \delta_{x_j,\Xi}\right)(f) &= \sum_{j=1}^{M} \alpha_j f(x_j) - \sum_{j=1}^{M} \alpha_j \left(\sum_{k=1}^{q} f(\xi_k) p_k(x_j)\right) \\ &= \sum_{j=1}^{M} \alpha_j f(x_j) - \sum_{k=1}^{q} f(\xi_k) \left(\sum_{j=1}^{M} \alpha_j p_k(x_j)\right)\end{aligned}$$

necessarily vanishes on $\mathcal{P}$ and is in $\mathcal{P}_\Omega^\perp$. Applying the conditional positive definiteness of Φ for this functional yields positivity of

$$\alpha^T A_{X,\Psi} \alpha$$

unless the coefficients of the above functional are zero, which implies that α is zero. □

By some simple linear algebra techniques the above normalization method can be shown to be equivalent to the method described in 8.2 (*Red2)* on page 109. To see this we give some hints, but suppress details of the full argument. Starting in section 8.2 (*Red2)* with a $\mathcal{P} = I\!P_m^d$-nondegenerate set $X = \{x_1, \ldots, x_M\}$, we can renumber the points and assume that $\Xi = \{x_1 \ldots, x_q\} = \{\xi_1 \ldots, \xi_q\}$ holds. Furthermore, if we pick the right basis in $I\!P_m^d$, the matrix S in (8.2.1, *Dec2*) has the elements $p_j(x_k),\ k = q+1, \ldots, M$. But then the matrix occurring in (8.2.4, *RedSys3*) precisely describes how to form the elements $\Psi(x_j, x_k)$ for $j, k = q+1, \ldots, M$ via the normalization formula (3.3.20, *DefSymmPsi*).

The function Ψ vanishes whenever one of its arguments is in Ξ. This is reflected in the above argument, since Ψ is responsible for reconstruction on $X \setminus \Xi = \{x_{q+1}, \ldots, x_M\}$.

3.3.3 Characterization of Native Spaces

(*SecCNS*) The native space associated to each conditionally positive definite function Φ is a rather abstract object, and it would be nice to know precisely which functions are in the space and which are not. This is a nontrivial task, since the only available information to start with is the conditional positive definiteness of Φ. Using transforms, we can give some results in section 4.2 (*SecCNST)*, and we leave the general case as an open research problem.

3.3.4 Remarks

The association of a Hilbert space to each conditionally positive definite function dates back to Madych and Nelson ([14](*madych-nelson:83-1*) [15](*madych-nelson:88-1*) [16](*madych-nelson:89-1*) [17](*madych-nelson:90-1*)).

3.3.5 Standardized Notation

(*SecSN*) The previous sections showed that it does not matter whether we start our theory from optimal recovery in spaces of functions with a bilinear form or from any given conditionally positive definite function. The only difference was that in the first case we constructed a normalized conditionally positive definite function from the given bilinear form, while in the second the given conditionally positive definite function Φ may not be normalized, though its normalization Ψ will generate the same bilinear form as Φ. From

now on we want to be independent from the starting point, and thus we collect the following facts that hold in both cases:

1. $\Phi : \Omega \times \Omega \to I\!R$ is a conditionally positive definite function on some domain Ω with respect to some nullspace $\mathcal{P}$ of finite dimension q.

2. There is a positive semidefinite bilinear form $(\cdot,\cdot)_\Phi$ on a space $\mathcal{G}$ of functions on Ω with nullspace $\mathcal{P}$.

3. The nullspace $\mathcal{P}$ has a basis $p_1,\ldots,p_q$ such that with certain linear functionals $\pi_1,\ldots,\pi_q$ on $\mathcal{G}$ the projector $\Pi_\mathcal{P}$ from $\mathcal{G}$ onto $\mathcal{P}$ is well-defined via
$$\Pi_\mathcal{P} g = \sum_{k=1}^{q} \pi_k(g) p_k \text{ for all } g \in \mathcal{G}.$$

4. For each $x \in \Omega$ the linear functionals
$$\delta_{x,\mathcal{P}} : g \mapsto g(x) - (\Pi_\mathcal{P} g)(x)$$
are continuous with respect to $(\cdot,\cdot)_\Phi$ and the Taylor-type reconstruction formula
$$\delta_{x,\mathcal{P}}(g) = (\delta^y_{x,\mathcal{P}} \Phi(y,\cdot), g(\cdot))_\Phi$$
holds for all $g \in \mathcal{G}$, $x \in \Omega$.

5. The space $\mathcal{G}$ can be decomposed into a direct sum $\mathcal{G} = \mathcal{P} + \mathcal{F}$ such that $\mathcal{F}$ is a Hilbert space with inner product $(\cdot,\cdot)_\Phi$.

6. If functionals $\lambda_{X,M,\alpha}$ are defined as
$$\lambda_{X,M,\alpha} : f \mapsto \sum_{j=1}^{M} \alpha_j f(x_j)$$
for sets $X = \{x_1,\ldots,x_M\} \subset \Omega$ and vectors $\alpha \in I\!R^M$ for arbitrary values of $M \geq q$, then one can define an inner product
$$(\lambda_{X,M,\alpha}, \lambda_{Y,N,\beta})_\Phi := \sum_{j=1}^{M}\sum_{k=1}^{N} \alpha_j \beta_k \Phi(x_j, y_k)$$
on all such functionals that vanish on $\mathcal{P}$. The set $\mathcal{P}^\perp_\Phi$ of all of these functionals then is an inner product space.

7. The space $\mathcal{G}$ is the largest space of functions on Ω such that all functionals in $\mathcal{P}^\perp_\Phi$ are continuous with respect to the norm induced by $(\cdot,\cdot)_\Phi$ on $\mathcal{P}^\perp_\Phi$.

8. The closure of $\mathcal{P}_\Phi^\perp$ under $(\cdot,\cdot)_\Phi$ is the dual $\mathcal{F}^*$ of $\mathcal{F}$, and the map $F : \mathcal{F}^* \to \mathcal{F}$ provides the Riesz correspondence between functionals and functions.

9. The action of F is related to Φ via

$$F(\lambda)(\cdot) = \lambda^y \Phi(y,\cdot)$$

for all $\lambda \in \mathcal{F}^*$. This is evident in case of $\lambda = \lambda_{X,M,\alpha} \in \mathcal{P}_\Phi^\perp$ and has to be read as a definition of the right-hand side for general λ.

10. The dual $\mathcal{G}^*$ of $\mathcal{G}$ consists of functionals that are sums of a linear functional on $\mathcal{P}$ and a linear functional in $\mathcal{F}^*$.

11. For each $\lambda \in \mathcal{G}^*$ we have $\lambda - \lambda\Pi_\mathcal{P} \in \mathcal{F}^*$ and

$$\begin{aligned}\lambda(g) &= \lambda\Pi_\mathcal{P} g + (F(\lambda - \lambda\Pi_\mathcal{P}), g)_\Phi \\ &= \lambda\Pi_\mathcal{P} g + (\lambda - \lambda\Pi_\mathcal{P}, F^{-1}(g - \Pi_\mathcal{P} g))_\Phi \\ &= \lambda\Pi_\mathcal{P} g + ((\lambda - \lambda\Pi_\mathcal{P})^y \Phi(y,\cdot), g)_\Phi\end{aligned}$$

for all $g \in \mathcal{G}$.

12. The normalization $\Psi_\mathcal{P}$ of Φ is defined via

$$\Psi_\mathcal{P}(x,y) = (\delta_{x,\mathcal{P}}, \delta_{y,\mathcal{P}})_\Phi = \delta_{x,\mathcal{P}}^u \delta_{y,\mathcal{P}}^v \Phi(u,v)$$

for all $x, y \in \Omega$. It has the properties described in Theorem 3.2.12 (*CPDFT1)*.

3.4 Error Bounds in Native Spaces of Functions

(*SecError*) This section starts our analysis of the error of optimal recovery processes. We first introduce the power functions of linear recovery processes in spaces of functions. It turns out that the power functions of optimal recoveries are pointwise optimal along all other power functions. This is the major tool for proving error bounds later, but we make a short detour concerning the stability of the recovery process and prove the Uncertainty relation in a general form.

3.4.1 Power functions

(*SecPF*) Assume that we have a quite general process that associates to each function g in a space $\mathcal{G}$ of functions on Ω another function $S(g) \in \mathcal{G}$ such that the map $S : g \mapsto S(g)$ is linear. The space $\mathcal{G}$ should carry at least a seminorm $|\cdot|$ with nullspace $\mathcal{P}$.

Definition 3.4.1 *The function*

(DefPowfct)

$$P(x) := P_{S,\mathcal{G}}(x) := \sup_{\substack{g \in \mathcal{G} \\ |g| \neq 0}} \frac{|(g - S(g))(x)|}{|g|} \in I\!R \cup \{\infty\} \tag{3.4.2}$$

is the **power function** *of S with respect to Φ.*

This is nothing else than the norm of the pointwise error functional if the latter is finite:

$$P(x) := \|\delta_{x,S}\| \text{ with } \delta_{x,S}(g) := g(x) - S(g)(x).$$

It yields the elementary error bound

(EqgSg)

$$|g(x) - S(g)(x)| \leq P(x)|g|, \;\; g \in \mathcal{G}, \; x \in \Omega. \tag{3.4.3}$$

If the projection property $S \circ S = S$ holds, then one can insert $g - S(g)$ instead of g into this bound to get

(EqgSg1)

$$|g(x) - S(g)(x)| \leq P(x)|g - S(g)|, \;\; g \in \mathcal{G}, \; x \in \Omega, \tag{3.4.4}$$

which often is some improvement over (3.4.3, *EqgSg*), because we frequently have $|g - S(g)| \leq |g|$.

To make the reader somewhat more familiar with the notion of a power function, we recall interpolation by univariate polynomials of order at most n on n distinct points $x_1 < \ldots < x_n$ in $[a, b] \subset I\!R$. The space $\mathcal{G}$ is $C^n[a, b]$ with seminorm $|g|_n := \|g^{(n)}\|_{[a,b],\infty}$, and the interpolant to g will be denoted by $S(g)$. The usual error bound

$$|g(x) - S(g)(x)| \leq \frac{1}{n!} |\prod_{j=1}^{n} (x - x_j)| |g|_n$$

is precisely of the form (3.4.3, *EqgSg*), and the power function is

$$P(x) = \frac{1}{n!} \prod_{j=1}^{n} |x - x_j|,$$

since it is well-known that the error bound is exact.

Power functions can be associated to almost every process of approximation or interpolation, and they enable comparison between different processes S on the same space $\mathcal{G}$ as well as the comparison of the same process S on different spaces $\mathcal{G}$, respectively.

3.4.2 Representations of Power Functions

(*SecRPF*) We now want to specialize the notion of a power function to the context of optimal recovery in function spaces. We assume the situation of Theorem 3.1.19 (*ORT2)* on page 34. That is, there are M linearly independent functionals $\lambda_1, \ldots, \lambda_M$ from $\mathcal{G}^*$ and a unique solution g^* of the optimal recovery problem (3.1.4, *ORP*). But we want to compare g^* with arbitrary other recoveries of g by linear methods that use the information $\lambda_j(g),\ 1 \le j \le M$. These have the form

(*GenRec*)

$$S_u(g) = \sum_{j=1}^{M} \lambda_j(g) u_j, \tag{3.4.5}$$

and we assume them to reproduce functions from $\mathcal{P}$. Then for each $x \in \Omega$ there is a functional

$$\delta_{x,u,S} \ : \ g \mapsto g(x) - S_u(g)(x) = g(x) - \sum_{j=1}^{M} \lambda_j(g) u_j(x)$$

vanishing on $\mathcal{P}$. The power function with respect to S_u is then representable via

$$P^2_{S_u,\Phi}(x) = |\delta_{x,u,S}|^2_\Phi.$$

It is now fairly easy to form

$$F(\delta_{x,u,S})(\cdot) = \Phi(x,\cdot) - \sum_{j=1}^{M} \lambda^z_j \Phi(z,\cdot) u_j(x)$$

and the function

(*DefPuxy*)

$$\begin{aligned} P_u(x,y) &:= (\delta_{x,u,S}, \delta_{y,u,S})_\Phi = \delta_{y,u,S}(F\delta_{x,u,S}) \\ &= \Phi(x,y) - \sum_{j=1}^{M} \lambda^z_j \Phi(z,y) u_j(x) \\ &\quad - \sum_{k=1}^{M} \lambda^z_k \Phi(,\cdot,z) u_k(y) + \sum_{j,k=1}^{M} \lambda^z_j \lambda^u_k \Phi(z,u) u_j(x) u_k(y) \end{aligned} \tag{3.4.6}$$

for all $x, y \in \Omega$. The reader will suspect some misuse of notation here, but the function $P_u(x,y)$ has some nice properties that justify this:

Theorem 3.4.7 *(PuT) The function $P_u(\cdot,\cdot)$ defined in (3.4.6, DefPuxy) satsfies*

- $P_u^2(x) = P_u(x,x) = \|P_u(x,\cdot)\|_\Phi^2$ *for all* $x \in \Omega$,
- $P_u(x,y) \le P_u(x)P_u(y)$ *for all* $x, y \in \Omega$,
- $P_u(x,\cdot)/P_u(\cdot)$ *attains its maximum* $P_u(x)$ *in* Ω *at* x,
- *if* $X = \{x_1, \ldots, x_M\} = \Xi = \{\xi_1, \ldots, \xi_q\}$ *is* $\mathcal{P}$*-nondegenerate and minimal, then* P_u *coincides with the normalization of* Φ *with respect to* Ξ,
- P_u *is another conditionally positive definite function that generates the same native space as* Φ.

Proof: The property $P_u^2(x) = P_u(x,x)$ follows from the definitions of both functions, and

$$\|P_u(x,\cdot)\|_\Phi = \|F(\delta_{x,u,S})\|_\Phi = \|\delta_{x,u,S}\|_\Phi$$

implies $\|P_u(x,\cdot)\|_\Phi = P_u(x)$. The next assertion is a consequence of the Cauchy-Schwarz inequality applied to the definition of $P_u(x,y)$, and together with the first it yields the third. The proof of the final property listed above is the same as for the normalization. □

The merit of (3.4.6, *DefPuxy*) is that it allows to write down the power function in explicit form and under quite general circumstances. This is of paramount importance for deriving error bounds in subsequent sections, and the basic feature is the optimality principle described in the next section.

3.4.3 Optimality of Power Functions of Optimal Recoveries

Equation (3.4.6, *DefPuxy*) defines $P_u(x,x) = P_u^2(x)$ for fixed x as a quadratic form of the M real variables $u_j(x)$, $1 \le j \le M$. We now want to minimize this quadratic form with respect to these variables, but we have to consider the restrictions

(*PolRepEq*)

$$\delta_{x,u,S}(p_i) = p_i(x) - \sum_{j=1}^{M} \lambda_j(p_i) u_j(x) = 0, \ 1 \le i \le q \tag{3.4.8}$$

imposed by reproduction of $\mathcal{P}$. Since $P_u^2(x)$ is nonnegative, the minimization must have a solution, and this solution can be characterized by the usual necessary conditions for quadratic optimization under linear constraints. There

must be Lagrange multipliers $\beta_1(x), \ldots, \beta_q(x)$ such that the solution $u_j^*(x)$ of the restricted optimization is a minimum of the unrestricted function

$$P_u^2(x) + \sum_{i=1}^{q} \beta_i(x) \left(p_i(x) - \sum_{j=1}^{M} \lambda_j(p_i) u_j(x) \right) = 0$$

of $u_1(x), \ldots, u_M(x)$. Taking the derivative with respect to $u_k(x)$, we get

$$0 = -2\lambda_k^z \Phi(z,x) + 2\sum_{j=1}^{M} \lambda_j^z \lambda_k^u \Phi(z,u) u_j^*(x) - \sum_{i=1}^{q} \beta_i(x) \lambda_k(p_i).$$

We can rewrite this together with (3.4.8, *PolRepEq*) to get the system

$$\begin{array}{rcrcll} \sum_{j=1}^{M} \lambda_j^z \lambda_k^u \Phi(z,u) u_j^*(x) & + & \sum_{i=1}^{q} \frac{-\beta_i(x)}{2} \lambda_k(p_i) & = & \lambda_k^z \Phi(z,x), & 1 \le k \le M \\ \sum_{j=1}^{M} \lambda_j(p_i) u_j^*(x) & + & 0 & = & p_i(x), & 1 \le i \le q \end{array}$$

The coefficient matrix of this system is the same as in (3.1.14, *EQsys3*) on page 32, if we use (3.2.14, *gjkrep*) on page 44. Thus the solution is in the span of the right-hand side, proving that $u_j^*(x) \in \mathcal{S},\ 1 \le j \le M$, as functions of x, but note that the necessary restriction on the $\beta_i(\cdot)$ of (3.1.33, *DefS*) is not satisfied. If we apply λ_k^x to these equations, we see that the conditions

$$\lambda_k^x(u_j(x)) = \delta_{jk},\ 1 \le j,k \le M$$

of interpolation are satsfied together with

$$\lambda_k^x(\beta_i(x)) = 0,\ 1 \le k \le M,\ 1 \le i \le q.$$

Thus we have

Theorem 3.4.9 *(OPFT) The power function $P_{u^*}(x)$ of the optimal recovery problem (3.1.4, ORP) is optimal with respect to u under all power functions $P_u(x)$ of recoveries of the form (3.4.5, GenRec) that reproduce $\mathcal{P}$.* □

This is in line with the optimality of the generalized optimal power function $P_\Lambda(\mu)$ of (3.1.38, *GPDef*) on page 39. Note that there we used optimal recovery right from the start, but allowed a general functional μ instead of a point evaluation functional δ_x, while in this section we allowed general recoveries, but restricted ourselves to the special functional δ_x. The explicit correspondence is

$$P_\Lambda(\delta_x) = P_{u^*}(x)$$

between these two versions of optimal power functions.

3.4.4 Condition

(*SecCondition*) We now look at the stability of solutions of the systems (1.7.2, *EQsys2*) and (3.1.14, *EQsys3*) written in matrix form as

(*BDef3*)

$$\begin{pmatrix} A & P \\ P^T & 0 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}. \tag{3.4.10}$$

which is exactly (1.7.3, *BDef*) or (3.1.18, *BDef2*), but repeated here for convenience. Introducing perturbations of the solution and the right-hand side we get the system

$$\begin{pmatrix} A & P \\ P^T & 0 \end{pmatrix} \begin{pmatrix} \alpha + \Delta\alpha \\ \beta + \Delta\beta \end{pmatrix} = \begin{pmatrix} f + \Delta f \\ 0 \end{pmatrix}$$

and can subtract (3.4.10, *BDef3*) to get

$$\begin{pmatrix} A & P \\ P^T & 0 \end{pmatrix} \begin{pmatrix} \Delta\alpha \\ \Delta\beta \end{pmatrix} = \begin{pmatrix} \Delta f \\ 0 \end{pmatrix}.$$

This implies

(*DAD*)

$$(\Delta\alpha)^T A(\Delta\alpha) = (\Delta\alpha)^T \Delta f. \tag{3.4.11}$$

Since we have $P^T(\Delta\alpha) = 0$, we know that the above quadratic form is positive definite. Thus there are positive real eigenvalues σ and Σ of the matrix A such that

$$0 < \sigma := \inf \frac{\alpha^T A\alpha}{\alpha^T\alpha} \leq \sup \frac{\alpha^T A\alpha}{\alpha^T\alpha} =: \Sigma < \infty,$$

where the sup and inf are extended over all $\alpha \in I\!R^M$ which are nonzero but satisfy $P^T\alpha = 0$. The **condition number** $\kappa(A)$ of A in the Euclidean norm then is the quotient $\kappa(A) = \Sigma/\sigma$, and it appears in the bound

$$\frac{\|\Delta\alpha\|_2}{\|\alpha\|_2} \leq \kappa(A)\frac{\|\Delta f\|_2}{\|f\|_2}$$

that follows from (3.4.11, *DAD*)and the corresponding equation

$$\alpha^T A\alpha = \alpha^T f$$

for the unperturbed quantities. This bound holds for the relative error, while the absolute error is governed by

(*Stab*)

$$\|\Delta\alpha\|_2 \leq \frac{1}{\sigma}\|\Delta f\|_2. \tag{3.4.12}$$

Numerical experiments show that σ can indeed be extremely small, while Σ does not grow too wildly, at least not as wildly as $1/\sigma$. Later theoretical results will support these statements, and thus the study of σ or some lower bounds for it will be of great importance for any assessment of the numerical stability of systems like (3.4.10, *BDef3*).

3.4.5 Remarks

The technique for proving error bounds via power functions goes at least back to Golomb and Weinberger [12](*golomb-weinberger:59-1*) but probably further back to Peano, since the error evaluation of linear functionals by bounding their Peano kernels is very similar. The pointwise optimality principle of Theorem 3.4.9 (*OPFT)* was used by various authors and possibly dates back to Duchon [6](*duchon:76-1*).

3.4.6 Uncertainty Relation

(*URT*) It would be very desirable to have recovery methods with small errors and good stability. However, these two goals cannot be met at the same time, since there is a connection between them that implies bad stability whenever the a-priori error bound is very small.

Let us look at this connection in a fairly general way. If we try optimal recovery of a function $g \in \mathcal{G}$ from data $\lambda_j(g), 1 \leq j \leq M$ in the setting of section 3.1.1 (*subsecORP)* and bound the error by Theorem 3.1.36 (*ORTFA)* on page 39, then we have to study the generalized optimal power function $P(\mu)$ of (3.1.38, *GPDef*), whose square has the representation (3.1.29, *BAPN*). But this quantity can be written as a value of the quadratic form associated to the matrix

$$A_{\mu,\Lambda} = \begin{pmatrix} (\mu,\mu)_\Phi & (\mu,\lambda_1)_\Phi & \dots & (\mu,\lambda_M)_\Phi \\ (\lambda_1,\mu)_\Phi & (\lambda_1,\lambda_1)_\Phi & \dots & (\lambda_1,\lambda_M)_\Phi \\ \vdots & \vdots & & \vdots \\ (\lambda_M,\mu)_\Phi & (\lambda_M,\lambda_1)_\Phi & \dots & (\lambda_M,\lambda_M)_\Phi \end{pmatrix}$$

with the vector $(1, -\alpha_1(\mu), \dots, -\alpha_M(\mu))^T \in \mathbb{R}^{M+1}$. This yields a lower bound

(*UR1*)

$$P^2(\mu) \geq \sigma(A_{\mu,\Lambda}) \left(1 + \|\alpha(\mu)\|_2^2\right) \geq \sigma(A_{\mu,\Lambda}) \tag{3.4.13}$$

for the power function in terms of the smallest eigenvalue of the matrix. This relates the error analysis to the stability analysis and provides the background for various cases of the Uncertainty Relation. Furthermore, it sets the direction for further progress: we need upper bounds for the power function P and positive lower bounds for the smallest eigenvalue σ. But we should be aware that the two sides of (3.4.13, *UR1*) behave differently as functions of Λ: the right-hand side will vanish, but not the left-hand side, if two functionals from Λ come too close to each other.

3.4.7 The Lagrange Case

We now specialize to the setting of Theorem 3.4.9 (*OPFT)* on page 67 with $X = \{x_1, \ldots, x_M\} \subset \mathbb{R}^d$ and $\Lambda = \{\delta_{x_1}, \ldots, \delta_{x_M}\}$. Then we have the matrix

$$A_{x,X} = \begin{pmatrix} \Phi(x,x) & \Phi(x,x_1) & \ldots & \Phi(x,x_M) \\ \Phi(x_1,x) & \Phi(x_1,x_1) & \ldots & \Phi(x_1,x_M) \\ \vdots & \vdots & & \vdots \\ \Phi(x_M,x) & \Phi(x_M,x_1) & \ldots & \Phi(x_M,x_M) \end{pmatrix}$$

and the vector $(1, -u_1^*(x), \ldots, -u_M^*(x))^T \in \mathbb{R}^{M+1}$ and get the special form

(*UR2*)

$$P_{u^*}^2(x) = P_\Lambda^2(\delta_x) \geq \sigma(A_{x,X}) \left(1 + \sum_{j=1}^{M} |u_j^*(x)|^2\right) \geq \sigma(A_{x,X}) \tag{3.4.14}$$

of (3.4.13, *UR1*). Note that both sides are continuous functions of x and X (or Λ standing for X) that vanish whenever x tends to points in X.

We now can give some hints to the results that follow in later sections. The Uncertainty Relation in the form (3.4.14, *UR2*) suggests to bound P^2 from above and σ from below, in order to have both upper bounds on the attainable error and on the numerical stability, measured by $1/\sigma$ due to (3.4.12, *Stab*). We shall see in 3.5 (*SecUBOPF)* that upper bounds for P^2 take the form

(*FBound*)

$$P_{u^*}^2(x) \leq F(h_{X,\Omega}) \text{ for all } x \in \Omega \tag{3.4.15}$$

where F is a monotonic function of the fill distance $h_{X,\Omega}$ defined in (2.1.2, *DDDef*) on page 17. On the other hand, the lower bounds for σ in 5.4 (*SecLBE)* will be of the form

(*GBound*)

$$\sigma(A_X) \geq G(s_X) \text{ for all } X = \{x_1, \ldots, x_M\} \subset \Omega \tag{3.4.16}$$

with the separation distance s_X defined in (2.1.1, *SDDef*). For gridded data on $\epsilon\mathbb{Z}^d \cap \Omega$ we can roughly expect $h_{X,\Omega} = s_X\sqrt{d}$, and then the Uncertainty Relation necessarily implies

(*UR3*)

$$F(t\sqrt{d}) \geq G(t) \tag{3.4.17}$$

for all $t \geq 0$. This allows to check the quality of the bounds (3.4.15, *FBound*) and (3.4.16, *GBound*), since the lowest possible bounds F and the largest possible bounds G must necessarily satisfy (3.4.17, *UR3*) and are optimal, if they turn (3.4.17, *UR3*) into an equality. This opens the race for optimal bounds of the form (3.4.15, *FBound*) and (3.4.16, *GBound*), and this text will describe the current state-of-the-art.

3.4.8 Remarks

The Uncertainty Relation seems to occur first in [28](*schaback:95-1*).

3.5 Upper Bounds for the Optimal Power Function

(*SecUBOPF*) Here we proceed to prove upper bounds of the form (3.4.15, *FBound*) for the optimal power function of optimal recovery. This approach uses results from classical approximation theory and does not require Fourier transforms. Another proof technique, using transforms, will follow in section 4.4 (*SecEBTrans*).

3.5.1 Assumptions and First Results

We specialize here to the case of multivariate Lagrange interpolation by conditionally positive definite functions $\Phi : \Omega \times \Omega \to \mathbb{R}$ of order m on some domain Ω that can be embedded into $\mathbb{R}^d$. The data locations are supposed to form a $\mathbb{P}_m^d$-nondegenerate set $X = \{x_1, \ldots, x_M\} \subset \Omega$, and we use functions u_j on Ω with (3.4.5, *GenRec*) that reproduce $\mathbb{P}_m^d$.

The power function with respect to these data and the functions u_j takes the special form

(*DefPuxyLag*)

$$\begin{aligned} P_u(x)^2 &:= \Phi(x,x) - 2\sum_{j=1}^{M} \Phi(x,x_j)u_j(x) \\ &\quad + \sum_{j,k=1}^{M} \Phi(x_j,x_k)u_j(x)u_k(x) \end{aligned} \tag{3.5.1}$$

from (3.4.6, *DefPuxy*). Note that we allow quite arbitrary u_j here in view of Theorem 3.4.9 (*OPFT*). If optimal recovery leads to Lagrange basis functions $u_j^*,\ 1 \le j \le M$, then

$$P_{u^*}(x) \le P_u(x)$$

holds for all $x \in \Omega$, yielding a pointwise upper bound for the optimal power function.

To start with, we fix a polynomial order $\ell \ge m$ and a point $x \in \Omega$. Around x we shall approximate Φ by a polynomial φ in the following sense:

Assumption 3.5.2 *(FBAss1) For each $x \in \Omega$ and a specific choice of a polynomial order ℓ there are positive constants ρ, h_0, and C_1 and a polynomial $\varphi : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ of order not exceeding ℓ in each d-variate variable, such that*

(PhiApp)

$$\|\Phi(x+u,x+v) - \varphi(u,v)\|_\infty \le C_1 h^\rho \tag{3.5.3}$$

for all $h \in [0,h_0]$ and all $u,v \in [0,h]^d$.

We shall vary x and ℓ later, and then all of the above quantities will be studied as functions of x and ℓ. Equation (3.5.3, *PhiApp*) may be viewed as resulting from a Taylor expansion around (x,x) or by an L_∞ approximation process. It is no drawback to assume symmetry of φ in the sense $\varphi(x,y) = \varphi(y,x)$, because the arithmetic mean of these two polynomials will do the job.

We now define a function Q_u^2 that serves as a polynomial approximation to P_u^2 near x, but which will turn out to be zero later:

(*DefQuxyLag*)

$$\begin{aligned} Q_u(x)^2 &:= \varphi(0,0) - 2\sum_{j=1}^{M} \varphi(0,x_j-x)u_j(x) \\ &\quad + \sum_{j,k=1}^{M} \varphi(x_j-x,x_k-x)u_j(x)u_k(x). \end{aligned} \tag{3.5.4}$$

Now it is time to specify our choice of $u_j,\ 1 \le j \le M$ via local polynomial reproduction of order ℓ near x. Since the dependence on x and h is crucial here, we stick to an explicit notation:

Assumption 3.5.5 *(FBAss2) For some $x \in \Omega$ and some $h \in [0, h_0]$ there is a subset $J_x(h)$ of $\{1, \dots, M\}$, positive constants $C_2(x,h), C_3(x,h)$, and a choice of M real numbers $u_1^h(x), \dots, u_M^h(x)$ such that*

(uDefJx)

$$J_x(h) \subseteq \{\ j\ :\ 1 \le j \le M,\ \|x - x_j\|_\infty \le C_2(x,h)h\}, \tag{3.5.6}$$

(uDef1)

$$u_j^h(x) = 0 \text{ for all } j \notin J_x(h), \tag{3.5.7}$$

(uDef2)

$$p(x) = \sum_{j \in J_x(h)} u_j^h(x) p(x_j) \text{ for all } p \in I\!P_\ell^d, \tag{3.5.8}$$

(uDef3)

$$1 + \sum_{j \in J_x(h)} |u_j^h(x)| \le C_3(x,h). \tag{3.5.9}$$

The first three items specify the local polynomial reproduction, while the last defines C_3 to be the corresponding Lebesgue constant. We apply (3.5.8, *uDef2*) to $\varphi(0, y - x)$ as a function of y to get

$$\varphi(0, x - x) = \varphi(0,0) = \sum_{j \in J_x(h)} u_j^h(x) \varphi(0, x_j - x)$$

to prove that Q_u is identically zero:

$$\begin{aligned} Q_u(x)^2 &= \varphi(0,0) - 2\varphi(0,0) + \sum_{j=1}^{M} \varphi(x_j - x, 0) u_j(x) \\ &= 0. \end{aligned}$$

We now bound the optimal power function by

(*FundBound*)

$$\begin{aligned} P^2_{u^*}(x) &\le P^2_u(x) \\ &= P^2_u(x) - Q^2_u(x) \\ &= \Phi(x,x) - \varphi(0,0) \\ &\quad -2\sum_{j\in J_x(h)} u^h_j(x)\left(\Phi(x,x_j) - \varphi(0,x_j-x)\right) \\ &\quad +\sum_{j,k\in J_x(h)} u^h_j(x)u^h_k(x)\left(\Phi(x_j,x_k) - \varphi(x_j-x,x_k-x)\right) \\ &\le \left(1+\sum_{j\in J_x(h)} |u^h_j(x)|\right)^2 C_1(x)(C_2(x,h)h)^\rho \\ &\le C_3(x,h)^2 C_1(x) C_2^\rho(x,h) h^\rho \end{aligned} \tag{3.5.10}$$

for all h with $C_2(x,h)h \le h_0$, where we have to keep in mind that everything still depends on ℓ. Nevertheless (3.5.10, *FundBound*) is the fundamental error bound for optimal power functions, and it can be applied in a large number of cases. We summarize:

Theorem 3.5.11 *(FundBoundT) Under the assumptions 3.5.2 (FBAss1) and 3.5.5 (FBAss2) the optimal power function has a local bound of order $\rho/2$ in x with respect to $h \to 0$, if the constants $C_2(x,h)$, $C_3(x,h)$ are bounded for $h \to 0$.* □

The applications of Theorem 3.5.11 (*FundBoundT)* come in two variations:

- To prove a **fixed** error order ρ, one fixes an appropriate ℓ and uses compactness arguments to bound all relevant "constants" with respect to x and h.
- To prove very strong non-polynomial error bounds like e^{-c/h^2} for fixed-scale Gaussians, one has to let ℓ tend to ∞ and study the variation of the "constants" with ℓ. This is a much harder task.

The two assumptions 3.5.2 (*FBAss1)* and 3.5.5 (*FBAss2)* require two different kinds of results to be proven in the following sections:

- an error bound for local polynomial approximation of Φ,
- and bounds on the Lebesgue constant for local polynomial interpolation in Ω.

3.5.2 Approximation Error in the Radial Case

(*SecAERC*) Here we consider the special situation of d-variate radial functions $\Phi(x,y) = \phi(\|x-y\|_2)$, and we want to check Assumption 3.5.2 (*FBAss1)*. The crucial term in (3.5.3, *PhiApp*) takes the form $\Phi(x+u, x+v) = \phi(\|u-v\|_2)$ and usually will not be nicely expandable into a polynomial in u and v. Fortunately, it is independent of x, since we are in a translation-invariant situation, and we only need an approximation to ϕ near zero. More precisely, we approximate $\phi(r)$ by a polynomial $p_n \in I\!P_n^1$ in r^2 on the domain $[0,h]$ for small $h > 0$ and define the error as

(*EDef1*)

$$\begin{aligned} E_n(\phi,h) &:= \inf_{p\in I\!P_n^1} \|\phi(r) - p(r^2)\|_{\infty,[0,h]} \\ &= \inf_{p\in I\!P_n^1} \|\phi(\sqrt{r}) - p(r)\|_{\infty,[0,h^2]}. \end{aligned} \tag{3.5.12}$$

This error can be bounded by univariate Jackson type theorems from classical approximation theory. Less sophisticated bounds simply take p as the Taylor expansion of $\phi(\sqrt{\cdot})$ in zero. With $\ell = 2n-1$ and h replaced by $2\sqrt{d}h$ one can put the result into Assumption 3.5.2 (*FBAss1)*.

Let us evaluate a few cases by standard techniques and cite the stronger Jackson results from the literature later.

Example 3.5.13 *(AEPHS) In the polyharmonic spline case $\phi(r) = r^\beta$ with $\beta \in I\!R_{>0}\backslash 2I\!N$ we can simply rescale the approximation problem to the interval $[0,1]$. That is,*

$$\begin{aligned} E_n(r^\beta,h) &= \inf_{p\in I\!P_n^1} \|r^{\beta/2} - p(r)\|_{\infty,[0,h^2]} \\ &= \inf_{p\in I\!P_n^1} \|(h^2 s)^{\beta/2} - p(h^2 s)\|_{\infty,[0,1]} \\ &= h^\beta \inf_{p\in I\!P_n^1} \|s^{\beta/2} - h^{-\beta} p(h^2 s)\|_{\infty,[0,1]} \\ &= h^\beta E_n(r^\beta,1). \end{aligned}$$

This yields the exact dependence on h and leaves the dependence on β to a classical Jackson result on $[0,1]$. We get $\rho = \beta$, and this is independent of $\ell = 2n-1$, provided that $\ell = 2n-1 \geq m \geq \lceil \beta/2 \rceil$ holds, since we have to exceed the order m of conditional positive definiteness. The most important cases $\beta = 1$ and $\beta = 3$ have the bounds $E_1(r,1) = 1/2$ and $E_2(r^3,1) = 2/27$, and these are available by direct analysis, using the Alternation Theorem of linear univariate L_∞ approximation. For this, see any textbook on Approximation Theory, e.g.: the classical books by Cheney [2](cheney:??-1) and Meinardus [3](meinardus:??-1).

Example 3.5.14 *(AETPS) Now consider the classsical thin-plate spline* $\phi(r) = r^\beta \log r$ *with* $\beta \in 2\mathbb{N}$ *and order* $m > \beta/2$ *of conditional positive definiteness. We proceed along the same lines and need* $\ell = 2n-1 \geq m > \beta/2$. *This implies* $\beta/2 \leq n-1$, *which is useful to get rid of the* log *term in*

$$\begin{aligned}
E_n(r^\beta \log r, h) &= \inf_{p\in\mathbb{P}_n^1} \|\frac{1}{2} r^{\beta/2}\log r - p(r)\|_{\infty,[0,h^2]} \\
&= \inf_{p\in\mathbb{P}_n^1} \|\frac{1}{2}(h^2 s)^{\beta/2}\log(h^2 s) - p(h^2 s)\|_{\infty,[0,1]} \\
&= \inf_{p\in\mathbb{P}_n^1} \|\frac{1}{2}(h^2 s)^{\beta/2}(\log(h^2)+\log s) - p(h^2 s)\|_{\infty,[0,1]} \\
&= \inf_{p\in\mathbb{P}_n^1} \|\frac{1}{2}(h^2 s)^{\beta/2}\log s - p(h^2 s)\|_{\infty,[0,1]} \\
&= \tfrac{1}{2}h^\beta \inf_{p\in\mathbb{P}_n^1} \|s^{\beta/2}\log s - h^{-\beta}p(h^2 s)\|_{\infty,[0,1]} \\
&= h^\beta E_n(r^\beta \log r, 1).
\end{aligned}$$

The case $\beta = 2$ *has* $E_2(r^2\log r, 1) = e^{-1}$.

Example 3.5.15 *(AEWF) Here we treat Wendland's [30](*wendland:95-1*) function* $\phi(r) = (1-r)_+^4(1+4r)$ *which is positive definite on* $\mathbb{R}^d$ *for* $d \leq 3$ *and in* $C^2(\mathbb{R}^d)$ *if* $r = \|x\|_2$ *for* $x \in \mathbb{R}^d$. *But our approach will be applicable to the whole class of piecewise polynomial functions of the form*

$$\phi(r) = \left\{ \begin{array}{rl} u(r^2) + r^{2n-1}v(r) & r \in [0,1] \\ 0 & \textit{else} \end{array} \right\},$$

where we pick a maximal n *such that* u *lies in* $\mathbb{P}_n^1$ *and* v *is an arbitrary univariate polynomial with* $v(0) \neq 0$. *This means that* u *covers the first terms of even degree, while* r^{2n-1} *is the first term of odd degree. This includes the full range of Wendland's functions from [30](*wendland:95-1*) as well as Wu's functions from [?](*wu:94-1*) for certain values of* n. *In case of* $\phi(r) = (1-r)_+^4(1+4r)$ *we have* $\phi(r) = 1 - 10r^2 + r^3(20 - 15r + 4r^2)$ *with* $n = 2$. *We now use* u *as an approximation to* ϕ *on small intervals. In particular,*

$$\begin{aligned}
E_n(\phi, h) &= \inf_{p\in\mathbb{P}_n^1} \|\phi(r) - p(r^2)\|_{\infty,[0,h]} \\
&\leq \|r^{2n-1}v(r)\|_{\infty,[0,h]} \\
&\leq C_5 h^{2n-1}
\end{aligned}$$

for $h \in [0,1]$ *with a suitable constant* C_5 *depending on* v, *e.g.:* $C_5 := \|v\|_{\infty,[0,1]}$. *Note that for the function* $\phi(r) = (1-r)_+^4(1+4r)$ *we get the same order as for the polyharmonic spline* $\phi(r) = r^3$.

So far we did not use sophisticated theorems from approximation theory, since we were interested in the correct power of h, not in the optimal behaviour of the bounds with respect to ℓ or n.

In the previous cases it did not make much sense to let ℓ or n be too large, because the approximation order with respect to h is not improved, and because we see later that large values of ℓ lead to bad Lebesgue constants when heading for Assumption (3.5.5, *FBAss2*). But the next case will be different in nature:

Example 3.5.16 *(AEGEl) The Gaussian $\phi(r) = \exp(-\alpha r^2)$ allows arbitrary values of $\ell = 2n-1$ because it is unconditionally positive definite. A crude bound is provided by chopping the exponential series:*

$$\begin{aligned} E_n(\exp(-\alpha r^2, h) &= \inf_{p \in I\!P^1_n} \|\exp(-\alpha r) - p(r)\|_{\infty,[0,h^2]} \\ &= \inf_{p \in I\!P^1_n} \|\exp(-s) - p(s/\alpha)\|_{\infty,[0,\alpha h^2]} \\ &\leq \frac{(\alpha h^2)^n}{n!} \end{aligned}$$

for $\alpha h^2 \leq n+1$, which is not a serious restriction. By using the Taylor residual one can get rid of the restriction, and by Bernstein's theorem 3.5.20 (BT1) on approximation of analytic functions we can get a similar bound that decays exponentially with $n \to \infty$. Anyway, we see that the bound improves dramatically with increasing n or $\ell = 2n-1$.

3.5.3 Jackson-Bernstein Theorems and Applications

This section contains the results from Approximation Theory that make the previous results somewhat sharper. We stick to radial functions and use univariate techniques. These consist of Jackson type theorems for the best approximation of functions $f \in C^n[a,b]$ by univariate polynomials in $I\!P^1_\ell$ in the supremum or Chebyshev or L_∞ norm:

(*EDef2*)

$$E(\ell, f, [a,b]) := \inf_{p \in I\!P^1_\ell} \|f - p\|_{\infty,[a,b]} \tag{3.5.17}$$

We additionally need the notion of **Lipschitz continuity**:

Definition 3.5.18 *A function f is Lipschitz continuous on $[a,b]$ of order $\alpha \in [0,1]$ with Lipschitz constant L, if*

$$|f(x) - f(y)| \leq L|x-y|^\alpha$$

holds for all $x, y \in [a,b]$.

Theorem 3.5.19 *(JT1) For all functions $f \in C^n[a,b]$ and all $\ell \geq n$ we have*

$$E(\ell, f, [a,b]) \leq \left(\frac{\pi}{4}\right)^n \frac{(b-a)^n}{(\ell+1)\ell \ldots (\ell-n+2)} \|f^{(n)}\|_\infty.$$

If $f^{(n)}$ is Lipschitz continuous of order α with Lipschitz constant L, then

$$E(\ell, f, [a,b]) \leq \left(\frac{\pi}{4}\right)^{n+1} \frac{(n+1)^n}{n!} \left(\frac{b-a}{\ell}\right)^{n+\alpha} L.$$

These results of Jackson (see e.g. Cheney [2](*cheney:??-1*) or Meinardus [3](*meinardus:??-1*)) yield bounds in terms of fixed negative powers of ℓ that depend on the smoothness of f. They can be proven to be optimal. For analytic functions, however, the parameter ℓ moves into the exponent of some quantity that is smaller than one, and this yields a much better asymptotic behaviour for $\ell \to \infty$ due to Bernstein (this is, for instance, in Natanson [5](*natanson:55-1*)) :

Theorem 3.5.20 *(BT1) Let f be a function on $[a,b]$ which has a holomorphic continuation into an ellipse in $\mathbb{C}$ with foci a, b and half-axes of length $0 < r \leq R$. Then there is a constant K depending only on f, r, and R, but not on ℓ, such that*

$$E(\ell, f, [a,b]) \leq K \left(\frac{b-a}{2(r+R)}\right)^\ell.$$

We cannot give proofs here, but the following weaker and easily accessible result shows how the previous result is possible.

Theorem 3.5.21 *(RSJT) Let f be a function on $[-r,r]$ which has a holomorphic continuation into the circle C_R in $\mathbb{C}$ with radius $R > r$ such that the continuation still is bounded on the boundary ∂C_R of the circle. Then*

$$E(\ell, f, [-r,r]) \leq \|f\|_{\infty, \partial C_R} \frac{R}{R-r} \left(\frac{r}{R}\right)^\ell,$$

and the bound is already achieved by the Taylor expansion around zero.

Proof: Just consider the power series of f in zero and bound it using Cauchy's inequality

$$|a_n| \leq \|f\|_{\infty, \partial C_R} R^{-n}$$

for the coefficients. This yields

$$\begin{aligned} |f(z) - \sum_{j=0}^{\ell-1} a_j z^j| &= |\sum_{j=\ell}^{\infty} a_j z^j| \\ &\le \sum_{j=\ell}^{\infty} |a_j| r^j \\ &\le \|f\|_{\infty,\partial C_R} \sum_{j=\ell}^{\infty} \left(\frac{r}{R}\right)^j \\ &\le \|f\|_{\infty,\partial C_R} \left(\frac{r}{R}\right)^{\ell} \frac{R}{R-r} \end{aligned}$$

for all $|z| \le r$. □

We now work our way through the examples.

Example 3.5.22 *(AEPHS2) Consider thin-plate splines $\phi(r) = r^\beta$. These are conditionally positive definite of order $m \ge m_\beta := \lceil \frac{\beta}{2} \rceil$. We have to approximate $r^{\beta/2}$ on $[0, h^2]$ and do this directly by application of Jackson's theorem 3.5.19 (JT1). The function $r^{\beta/2}$ has $m_\beta - 1$ continuous derivatives, and the final derivative is Lipschitz continuous of order*

$$\alpha_\beta := \frac{\beta}{2} - m_\beta + 1 = \frac{\beta}{2} - \lfloor \frac{\beta}{2} \rfloor \in (0, 1)$$

with constant

$$L_\beta = \frac{\beta}{2}(\frac{\beta}{2} - 1) \dots (\frac{\beta}{2} - m_\beta) = (1 + \alpha_\beta)(2 + \alpha_\beta) \dots (m_\beta - 1 + \alpha_\beta) \le m_\beta!.$$

Then the two slightly different notions of (3.5.12, EDef1) and (3.5.17, EDef2), which are related by the transformation $r \mapsto \sqrt{r}$ in the argument of the function, come out to be

$$E_n(r_\beta, h) = E(n, r^{\beta/2}, [0, h^2]) \le \left(\frac{\pi}{4}\right)^{m_\beta} \frac{(m_\beta)^{m_\beta - 1}}{(m_\beta - 1)!} \left(\frac{h^2}{n}\right)^{\beta/2} L_\beta$$

for all $\ell = 2n - 1 \ge m \ge m_\beta$. The result has the same power of h as before, but now we can quantify the dependence on β and n. Unfortunately, the gain for large n or $\ell = 2n - 1$ is much too weak to cope with the dramatic increase of Lebesgue constants for increasing polynomial degrees.

Example 3.5.23 *(AETPS2) We now continue with Example 3.5.14 (AETPS). The radial function $\phi(r) = r^\beta \log r$ with $\beta \in 2\mathbb{N}$ is conditionally positive definite of order $m \ge m_\beta := \beta/2 + 1$. We have to consider polynomial approximations to $r^{\beta/2} \log r$ for orders n satisfying $\ell = 2n - 1 \ge m \ge m_\beta = \beta/2 + 1$.*

The derivatives of $r^{\beta/2}\log r$ for $\beta \in 2\mathbb{N}$ produce lower-order polynomials of type $r^{\beta/2-1}, r^{\beta/2-2}, \ldots$ which are subsumed in $\mathbb{P}_n^1$ and do not change the approximation error. Thus we only have to consider the terms of type $r^\alpha \log r$, and we see that we can take $\beta/2-1$ continuous derivatives. The final derivative is $(\beta/2)!r\log r$, which is Lipschitz continuous of order < 1, but not of order 1. The direct application of the second version of Jackson's theorem 3.5.19 (JT1) would not give the full order with respect to h due to this fact, and therefore we first do the scaling of Example 3.5.14 (AETPS) to extract the factor h^β out of $E_n(r^\beta \log r, h)$. Then the first version of Jackson's theorem yields

$$\begin{aligned} E_n(r^\beta \log r, h) &= h^\beta E_n(r^\beta \log r, 1) \\ &= E(n, r^{\beta/2}\log r, [0,1]) \\ &\le \left(\tfrac{\pi}{4}\right)^{\beta/2} \frac{(\beta/2)!}{(n+1)n(n-1)\ldots(n-\beta/2+2)} \|r\log r\|_{\infty,[0,1]} \\ &= \left(\tfrac{\pi}{4}\right)^{\beta/2} \left(\begin{array}{c} n+1 \\ \beta/2 \end{array}\right)^{-1} e^{-1} \end{aligned}$$

for all $\ell = 2n-1 \ge m \ge m_\beta = \beta/2+1$. Again, we have some improvement for increasing n, but it will not be enough to cope with the Lebesgue constants.

Example 3.5.24 *(AEMQ) We now consider multiquadrics $\phi(r) = (c^2 + r^2)^{\beta/2}$ for $\beta \notin 2\mathbb{N}$ and $c > 0$. In case of $\beta > 0$ they are conditionally positive definite of order $m \ge m_\beta := \lceil \beta/2 \rceil$, while they are positive definite for $\beta < 0$. In this case we define $m_\beta := 0$. Multiquadrics are analytic around $r = 0$ and their polynomial approximation can be treated by application of Bernstein's theorem 3.5.20 (BT1) or by Theorem 3.5.21 (RSJT). This means that we should study the complex function $f(z) = (c^2+z)^{\beta/2}$ which has a singularity at $z = -c^2$. For $\beta > 0$ the function is bounded on the circle C_{c^2}, but for negative β (inverse multiquadrics) we have to use a smaller radius. To be safe, we use $R = c^2/2$ in both cases and get*

$$|f(z)| \le (3c^2/2)^{\beta/2} \le 2^{|\beta/2|} c^\beta$$

for $\beta > 0$ and $|z| = R$, while

$$|f(z)| \le (c^2/2)^{\beta/2} = 2^{|\beta/2|} c^\beta$$

for $\beta < 0$ yields the same bound. We approximate on $[0, h^2]$ and thus have the constraint

$$h^2 < R = c^2/2$$

on what follows. Now Theorem 3.5.21 (RSJT) yields

(MQB1)

$$E_n(\phi, h) = E(n, f, [-h^2, h^2]) \le 2^{|\beta/2|} c^\beta \frac{c^2}{c^2 - 2h^2} \left(\frac{2h^2}{c^2}\right)^n \tag{3.5.25}$$

for all $\ell = 2n - 1 \ge m \ge m_\beta$.

Example 3.5.26 *(AES1) We consider Sobolew radial basis functions*

$$\phi(r) = r^\nu K_\nu(r)$$

for $\nu > 0$. *These generate Sobolev spaces* $W_2^m(\mathbb{R}^d)$ *for* $\nu = m - d/2$ *and are unconditionally positive definite. A direct application of Jackson's theorems requires the derivatives of* ϕ, *which are not easy to calculate and bound from above. We postpone treatment of this case to section 4.4.1 (EBSob), where we apply Fourier transform techniques.*

3.5.4 Lebesgue Constants

(*SecLebCon*) We now face the verification of Assumption 3.5.5 (*FBAss2)*, which is a very hard problem. Let us first discuss some easy cases.

3.5.5 Lines and Triangles

Assume that we want to prove a bound for the error in a point x that lies on a line between two distinct data points, say x_1 and x_2, and assume that the distance between these points is $2h$. We can define linear functions u_1, u_2 by

$$u_1(y) := \frac{(y - x_2)^T (x_1 - x_2)}{\|x_1 - x_2\|_2^2},\ u_2(y) := 1 - u_1(y)$$

and see that $u_j(x_k) = \delta_{jk},\ j, k = 1, 2$. Any linear polynomial p restricted to the line through x_1 and x_2 is uniquely recovered by $p(x) = p(x_1)u_1(x) + p(x_2)u_2(x)$. Note that Assumption 3.5.5 (*FBAss2)* only requires the recovery in x, not everywhere. If x is way between x_1 and x_2, then clearly $C_3 = 2$ suffices, since both $u_1(x)$ and $u_2(x)$ are in $[0, 1]$ and sum up to 1. Furthermore, we can set $C_2 = 1$ and are done for cases with $\ell \le 2$. This argument works for every space dimension, but only on lines between two nearby data points.

We now go over to three points x_1, x_2, $x_3 \in \mathbb{R}^d$ forming a nondegenerate triangle T, and we consider points x inside such a triangle. If x lies on an edge, we are in the previous case. Since our argument is carried out in a two-dimensional affine subspace containing the triangle, we assume that we are in

R^2 right away, and there are no problems going back to the embedded plane in $\mathbb{R}^d$. Nondegeneracy of the triangle, when written in bivariate coordinates, means that the system

$$\begin{pmatrix} x_1 & x_2 & x_3 \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} u_1(y) \\ u_2(y) \\ u_3(y) \end{pmatrix} = \begin{pmatrix} y \\ 1 \end{pmatrix}$$

has a nonsingular matrix and a unique solution. The components of the solution are called the **barycentric coordinates** of y with respect to the triangle spanned by $x_1,\ x_2,\ x_3$, and they satisfy

- $u_j(y)$ is linear in y,
- $u_j(x_k) = \delta_{jk},\ 1 \le j,k \le 3$,
- $p(y) = \sum_{j=1}^{3} u_j(y)p(x_j)$ for all $p \in \mathbb{P}_2^2$,
- $u_j(y) = 0$ iff y lies on the boundary line opposite to x_j,
- all $u_j(y) > 0,\ 1 \le j \le 3$ iff y lies inside the triangle,
- the $u_j(y)$ are nonnegative and sum up to 1 for y not outside the triangle.

The reader will have noticed that this is a very simple generalization from the two-point case. This can be carried further, but it never yields more than reproduction by linear polynomials. It always works for $d+1$ points that lie at least in $\mathbb{R}^d$ but not in a $d-1$-dimensional affine subspace.

It is now clear that in our three-point case we get $C_3 = 2$ independent of x and h, and if h is taken as the fill distance 2.1.2 (*DDDef)*

$$h := h_{\{x_1,x_2,x_3\},T} := \sup_{x \in T} \min_{1 \le j \le 3} \|x - x_j\|_2,$$

of the triangle T, then we have $C_2 = 1$. This argument works on all small triangles that are formed by three data points that are not on a line.

We now assemble the two cases into a general strategy that works in $\mathbb{R}^2$ for polynomial reproduction of order $\ell \le 2$. Assume that the set $X = \{x_1, \ldots, x_M\} \subset \mathbb{R}^2$ of scattered data is given, and let Ω be the convex hull of X, i.e.: the smallest convex set containing X. Then Ω is a compact convex polygon, and each point x of Ω either lies on a line between two points of X

or in a nondegenerate triangle formed by three points of X. Assume that X fills Ω with a fill distance

$$h := h_{X,\Omega} := \sup_{x\in\Omega} \min_{1\le j\le M} \|x - x_j\|_2.$$

If the situation of one of the two above cases occurs, there will not necessarily be two points on a line with distance at most $2h$ or a triangle T with local fill distance h. We thus have to determine which distances as factors of h are possible in these cases. We form the Delaunay triangulation of the set $X = \{x_1, \ldots, x_M\}$ as described in section 9.1 (*SecVor)*. This splits Ω into triangles with vertices at the points of X, and where there is an edge from x_k to x_j iff the midpoint between x_k and x_j has both x_k and x_j as points of X with minimal distance. Since this distance is at most h, the Delaunay triangles have edges of length at most $2h$. If we work on a line joining two vertices of the Delaunay triangulation, we thus have $C_2 = 1$. Inside of such triangles, the maximum distance from an interior point to the vertices is achieved in the isosceles case, and thus the fill distance within Delaunay triangles is at most $2h/\sqrt{3}$. We thus get away with $C_2 = 2/\sqrt{3}$ and $C_3 = 1$ in both cases.

3.5.6 Univariate Data

The situation for local polynomial interpolation of order exceeding two is much harder, even in one space dimension, where the solution still can be given using elementary techniques. Let us do a simple, but nonoptimal bound. Consider an odd number $\ell = 2k + 1$ points ordered locally on the real line like

$$x_1 < x_2 < \ldots < x_\ell$$

and let the fill distance of $X = \{x_1, \ldots, x_\ell\}$ be $h/2$, such that we have $x_{i+1} - x_i \le h$. The Lagrange basis functions for interpolation of order ℓ are

$$u_j(x) = \prod_{i\ne j} \frac{x - x_i}{x_j - x_i},\ 1 \le j \le \ell$$

and they get large if there are points with $x_j - x_i$ extremely small. But those points can be neglected if our points are a local subsample of a much larger set. Let us thus assume a real number $\alpha \in (0, 1]$ with $\alpha h \le x_{i+1} - x_i \le h$. Then the $M - 1 = 2k$ factors in the numerator can be bounded above by $(2k)!h^{2k}$, while the denominator can be bounded below by $(k!)^2\alpha^{2k}h^{2k}$. We have to sum $M = 2k + 1$ of these quotients, and thus

$$C_3 \le 1 + \frac{(2k + 1)!}{(k!)^2\alpha^{2k}},\ \ell = 2k + 1.$$

This bound increases dramatically with ℓ, unfortunately, but it is independent of h. We can get an idea of the behaviour of C_3, if we apply **Stirling's inequality**

(*Stirling*)

$$1 \leq \frac{n!}{\sqrt{2\pi n} n^n e^{-n}} \leq \exp(-\frac{1}{12n}). \tag{3.5.27}$$

The result is

$$C_3 \leq \mathcal{O}\left(\sqrt{k}\left(\frac{2}{\alpha}\right)^{2k}\right)$$

for $k \to \infty$ or in simplified form and as a function of ℓ,

$$C_3 \leq \mathcal{O}(\gamma^\ell)$$

for $\ell \to \infty$ with some $\gamma > 1$.

Now let us apply this globally, and assume an ordered, but scattered set $X = \{x_1, \ldots, x_M\} \subset I\!R$ with fill distance h. For a uniform distribution of points of meshwidth $3h$ over $\Omega = [x_1, x_M]$ we associate a scattered point to each meshpoint, and then this selection of a subset of $N \leq M$ points has a fill distance of $2.5h$ and each adjacent pair of points is at least h and at most $5h$ apart. We then can apply the above bound with $\alpha = 1/5$ by local selection of ℓ points and an appropriate scaling. If we use a uniform distribution with spacing Kh, we end up with $\alpha = (K-2)/(K+2)$ which can be pushed towards 1 for K large. To check the value of C_2, we have to assume the worst case, in which some x lies at the boundary, while the next interpolation point is $2.5h$ away, and the interpolation points are at maximal distance $5h$. Then the maximal value of $|x - x_j|$ is $2.5h + (\ell - 1) \cdot 5 \cdot h < 5\ell h$ such that we have $C_2 = 5\ell$.

The above approach is unfeasible for multivariate cases, because we relied heavily on the ordering of the points. But it gives us two pieces of useful information: the good news is that we might get along with a quantity C_3 that does not depend on h, but the bad news is that C_3 will crucially depend on the order ℓ of local polynomial interpolation. We address the two topics one after another.

3.5.7 Independence of h

As we saw in the univariate case, one can expect that the scaling parameter h cancels out in the bounds for C_3. To generalize this statement, we repeat the technique that we already used before:

1. If a set $X = \{x_1, \ldots, x_M\}$ is given in Ω with fill distance h, we pick an integer $k \geq 3$ and lay a grid $G = kh\mathbb{Z}^d$ over Ω.

2. For each point from $G \cap \Omega$ we pick the nearest data point from X. This yields a subset Y of X of points that are only mildly scattered and are at least $(k-2)h$ apart from each other. We need this to avoid degeneration of the local polynomial interpolation that we want to construct. Since the diagonal in the unit cube in R^d has length $\sqrt{d}$, the fill distance of Y in Ω is at most $h(1 + k\sqrt{d})$.

3. If $x \in \Omega$ is given, we pick a selection of points from Y which are near to x, and the indices of these points define the set $J_x(h)$ occurring in Assumption 3.5.5 (*FBAss2)*.

4. The main problem now is to prove that the selection guarantees solvability of polynomial interpolation of order ℓ.

5. We then evaluate the Lebesgue constants for this local interpolation.

If k is large, the set Y will consist of points that are relatively near to the grid $G = kh\mathbb{Z}^d$, since they can be only h away from gridpoints. Thus the local interpolation takes place on data that are slight perturbations of gridded data. We thus have to study polynomial interpolation on gridded data first, and then ask for admissible perturbations.

We write multivariate polynomials $p \in \mathbb{P}_\ell^d$ as

(*PolRep*)

$$p(x) = \sum_{|\alpha|<\ell} p_\alpha x^\alpha \tag{3.5.28}$$

with the usual multiindex notation:

$$\alpha \in \mathbb{Z}_{\geq 0}^d,\ |\alpha| := \|\alpha\|_1,\ x^\alpha := \prod_{j=1}^d x_j^{\alpha_j}.$$

The number of data points should equal the number of basis functions, and thus we simply use the data set

$$X_\ell^d := \{\ \beta \in \mathbb{Z}_{\geq 0}^d\ :\ |\beta| < \ell\ \}.$$

For $d = 2$ these are the points $(j,k) \in \mathbb{Z}^2$ with $0 \leq j,k \leq j+k < \ell$ forming a "triangle" in $\mathbb{Z}^2$.

Lemma 3.5.29 *(LemPIG) The set X_ℓ^d is a minimal nondegenerate set in $\mathbb{R}^d$ for polynomials in $\mathbb{P}_\ell^d$. Thus polynomial interpolation of order ℓ is uniquely possible.*

Proof: Since the dimension of $\mathbb{P}_\ell^d$ and the number of points in X_ℓ^d agree, it suffices to prove nondegeneracy. Let p be a polynomial of the form (3.5.28, *PolRep*) that vanishes on X_ℓ^d, and we want to show that p vanishes everywhere. We do this by induction on the space dimension d, and the case $d = 1$ is well-known. So we assume that for $k < \ell$ all polynomials from $\mathbb{P}_k^d$ that vanish on X_k^d must be identically zero. Now we extract the variable x_d from each of the terms in (3.5.28, *PolRep*), split x as $x = (\tilde{x}, x_d)$, and rearrange the sum. This yields

$$p(x) = p(\tilde{x}, x_d) = \sum_{j=0}^{\ell-1} p_j(\tilde{x}) x_d^j$$

with polynomials $p_j \in \mathbb{P}_{\ell-j}^{d-1}$. Setting $x = (0, k)$ for $0 \le k < \ell$ we see that $(0, k) \in X_\ell^d$ and the univariate polynomial

$$p(0, x_d) = \sum_{j=0}^{\ell-1} p_j(0) x_d^j$$

in $\mathbb{P}_\ell^1$ vanishes in the ℓ distinct points $k,\ 0 \le k < \ell$. Thus it is zero as a polynomial in x_d, proving $p_j(0) = 0,\ 1 \le j < \ell$. Looking at $p_{\ell-1} \in \mathbb{P}_1^{d-1}$ we see that $p_{\ell-1}$ must be zero.

Now let us start an inner induction over an integer $j = 1, 2, \ldots$ and assume that we already have proven that $p_{\ell-j}, \ldots, p_{\ell-1}$ are identically zero, and that all of the p_i vanish on X_j^{d-1}. This is precisely what we have proven for $j = 1$ and $X_1^{d-1} = \{0\}$. Now fix an arbitrary $\beta \in X_{j+1}^{d-1} \setminus X_j^{d-1}$. Then $|\beta| = j$ and we can form the data points $x = (\beta, k) \in X_\ell^d$ for $0 \le k < \ell - j$. Considering these points, the univariate polynomial

$$p(\beta, x_d) = \sum_{i=0}^{\ell-j-1} p_i(\beta) x_d^i$$

must have vanishing coefficients, and thus all p_i vanish on X_{j+1}^{d-1}. For $p_{\ell-j-1} \in \mathbb{P}_{j+1}^{d-1}$ the hypothesis of the outer induction yields that this polynomial vanishes identically, and this finishes the inner and outer induction. □

Now we know that (in a fixed enumeration of X_ℓ^d used for rows as well as columns) the matrix with elements α^β for $\alpha, \beta \in X_\ell^d$ is nonsingular. It is a

continuous function of the data locations α, and thus it is still nonsingular when all the points vary in local balls of some positive radius $\rho \in (0, 1/2)$ around the integer points of X_ℓ^d. This radius is a function of both d and ℓ, and to give an explicit positive lower bound for it is a formidable task. We neglect this problem now and consider d and ℓ as fixed, leading to some mysterious, but clearly positive radius ρ for admissible perturbations.

Lemma 3.5.30 *(LemPIP) For each space dimension d and each polynomial order ℓ there is a positive quantity $\rho(d, \ell)$ such that interpolation by polynomials in $I\!P_\ell^d$ is uniquely possible in all data sets that pick a point in each of the balls*

$$B_\rho(\alpha) := \{\, y \in I\!R^d \;:\; \|y - \alpha\|_2 \le \rho\}$$

for all $\alpha \in X_\ell^d$. The maximum Lebesgue constant for all of these polynomial interpolation processes, measured on a fixed ball B_R of some large radius R containing the set X_ℓ^d is some finite positive quantity $C_3^*(d, \ell, R)$.

Proof: Each pick of points defines a nice interpolation problem that has Lagrange basis functions $\{u_\alpha\}_\alpha$ depending continuously on the locations of the points. Thus also the Lebesgue constant

$$1 + \sup_{x \in B_R} \sum_{\alpha \in X_\ell^d} |u_\alpha(x)|$$

varies (via the u_α) continuously with the data locations. Since these vary in a compact set, the Lebesgue constants, as defined above, attain a finite maximum under variation of the data locations. □

Of course, one could replace the domain B_R of "measurement" in the Lebesgue constant by any compact set in $I\!R^d$, but note that the actual upper bound of the Lebesgue constants remains mysterious, and enlarging B_R will have a nasty blow-up effect.

The next step is the independence of the above situation under shifts and scaling:

Lemma 3.5.31 *(LemPIS) Let $Y \subset I\!R^d$ be a data set where interpolation by $I\!P_\ell^d$ is uniquely possible, and let $\{u_y(\cdot)\}_{y \in Y}$ be the associated Lagrange basis satisfying $u_y(x) = \delta_{xy}$ for $x, y \in Y$. If Y is translated by some $z \in I\!R^d$ and scaled by some $h > 0$ to go over into*

$$Z := h(Y - z) := \{\, y_h := h(y - z) \;:\; y \in Y\},$$

then interpolation in Z is equally possible, using the basis functions

$$u_{y_h}(\cdot) := u_y(z + \cdot/h)$$

and yielding the same Lebesgue constants, if the domain of measurement of those is translated and scaled accordingly, i.e.: the domain B is transformed into $h(B-y)$.

Proof: The definition of the new functions makes sure that they are polynomials in $I\!P_\ell^d$ and satisfy the Lagrange interpolation property. Looking at the definition of the Lebesgue constant proves the rest. □

We now go back to our data set $X = \{x_1, \ldots, x_M\}$ that fills Ω with a fill distance h, and we pick points from X that are perturbations of points from a grid $kh\mathbb{Z}^d$ laid over Ω. The perturbations thus stay within h of the grid points, while these are kh apart along the axes. Scaling them down by division with kh will bring them to the unit grid Z^d, and the scaled perturbations will stay within a radius $1/k$. We thus have to make sure that $1/k \leq \rho(d,\ell)$ holds and that we use a shifted, scaled, and perturbed version of X_ℓ^d for local interpolation. The point x must lie in the shifted and scaled domain of measurement of the Lebesgue constant. We then can use the bound $C_3^*(d,\ell,R)$ of the Lebesgue constant from Lemma 3.5.30 (*LemPIP)* for all h, as asserted.

Thus we are left to determine the constant C_2 of Assumption 3.5.5 (*FBAss2)* that bounds the maximal distance of x in terms of multiples of h to the points we use for interpolation. This is no big deal when x is in the interior of Ω and h is small enough, and we can then get away with something like $C_2 = k\ell\sqrt{d}$, the diameter of the cube $[0, k\ell]^d$. If x lies near the boundary of Ω, we must be more careful, because the boundary could have awkward outgoing cusps that take boundary points far away from places where we can find enough data from X that lie near the gridpoints of $kh\mathbb{Z}^d$ and allow full interpolation up to order ℓ. We make the following assumption:

Definition 3.5.32 *(DefICC) A closed compact domain $\Omega \in I\!R^d$ with nonempty interior satisfies an* **interior cone condition***, if there is a fixed positive angle γ and a fixed height δ such that for any boundary point x there is a cylindrical cone within Ω that has vertex x, angle γ at the vertex, and height δ.*

If Ω satisfies an interior cone condition, we can consider coverings of Ω by fine grids $\epsilon\mathbb{Z}^d$, and we see that there is a constant K_c such that for all ϵ that are small enough, i.e. $\epsilon \leq \epsilon_c$, any point of Ω is only $K_c\epsilon$ away from a grid cell of $\epsilon\mathbb{Z}^d$ that is completely contained in Ω. We apply this for $\epsilon := kh\ell \leq \epsilon_c$ and get that any x is at most $K_c\epsilon = K_c kh\ell$ away from a fully interior cell of sidelength $kh\ell$ in which we can do the local interpolation.

Thus all interpolation points will be at most $(K_c+1)kh\ell\sqrt{d}$ away from x and we can use $C_2 = (K_c+1)k\ell\sqrt{d}$.

We have two restrictions up to now:

$$1/k \le \rho(d,\ell) \text{ and } kh\ell \le \epsilon_c.$$

This yields the conditions

$$h \le h_c := \frac{\epsilon_c}{k\ell} \text{ and } k \ge \frac{1}{\rho(d,\ell)}$$

which are no problem for fixed values of d, ℓ, and the interior cone condition on Ω. We summarize:

Theorem 3.5.33 *(LPIT) For given values of d and ℓ and for a fixed cone condition there are positive constants h_c, C_2, C_3 such that Assumption 3.5.5 (FBAss2) is valid for all x in domains $\Omega \subset \mathbb{R}^d$ satisfying the cone condition, and for all $h \le h_c$.* □

Theorem 3.5.33 (*LPIT*) is useful for all cases where the local approximation uses only a finite degree ℓ, and where the exact value of the constants does not matter much. We defer a more detailed analysis to the next lemmas, where we rely on [18](*madych-nelson:92-1*).

To treat the general case, we cite without proof a deep result from [18](*madych-nelson:92-1*):

Theorem 3.5.34 *(MNL) Let R be a cube in $\mathbb{R}^d$ which is divided into K^d identical subcubes for some large integer K, and define $\gamma_d := 2d(1+\gamma_{d-1})$ starting with $\gamma_1 := 2$. Consider arbitrary polynomials p from $\mathbb{P}_\ell^d$ and assume $K \ge \ell\gamma_d$. Then*

(MNLBound)

$$\|p\|_{\infty,R} \le e^{2d\ell\gamma_d}\|p\|_{\infty,Y} \tag{3.5.35}$$

holds for any set $Y \subset R$ that picks a point from each of the subcubes. In particular, all these sets Y are $\mathbb{P}_\ell^d$-nondegenerate.

We now bring this into line with Lemma 3.5.30 (*LemPIP*) and assume $\ell \ge 2$ throughout. We want to let the little subcubes be centered around the points of X_ℓ^d. If their sidelength is 2ρ to make balls of radius ρ safely contained in the cubes, we have to take $(2\rho)^{-1} =: M \in \mathbb{N}$ and let the large cube be $R = [-\rho, \ell-1+\rho]^d$. Splitting it into K^d subcubes yields the equation

$$2\rho = \frac{\ell - 1 + 2\rho}{K}$$

which leads to

$$K = 1 + M(\ell - 1),\ \rho = \frac{\ell - 1}{2(K-1)},\ M = \lfloor \frac{\ell\gamma_d - 1}{\ell - 1} \rfloor.$$

This allows the application of Theorem 3.5.34 (*MNL)*. We first check the value of ρ as a function of d and ℓ. If we bound M crudely from above by $2\gamma_d$, we get $\rho \geq (4\gamma_d)^{-1}$, which is independent of ℓ.

The linear functional $\delta_x\ :\ p \mapsto p(x)$ can be written in the form

$$\sum_{\alpha \in X_\ell^d} p(z_\alpha) u_\alpha(x)$$

where we have picked z_α from the ball $B_\rho(\alpha)$, and where the u_α are the Lagrange interpolation polynomials. Then we fix $x \in R$ and interpolate data sgn $(u_\alpha(x))$ in z_α by some polynomial $\tilde{p} \in I\!P_\ell^d$ and get the bound

$$\sum_{\alpha \in X_\ell^d} |u_\alpha(x)| = \tilde{p}(x) \leq \|\tilde{p}\|_{\infty,R} \leq e^{2d\ell\gamma_d}.$$

Thus the Lebesgue constant in Lemma 3.5.30 (*LemPIP)* is bounded by

(*C3def*)

$$C_3^* \leq 1 + e^{2d\ell\gamma_d}. \tag{3.5.36}$$

We now go over to the situation in Theorem 3.5.33 (*LPIT)*. We fix a cone condition and a space dimension. This fixes the constants $\epsilon_c,\ h_c$, and K_c from the cone condition. The integer k can be chosen as $k = 4\gamma_d$ to satisfy $k\rho \geq 1$, and we are left with the condition

(*hrestr*)

$$h \leq h_c = \frac{\epsilon_c}{4\ell\gamma_d} \tag{3.5.37}$$

under which we can use (3.5.36, *C3def*) and

(*C2def*)

$$C_2 = 4(K_c + 1)\ell\sqrt{d}\gamma_d. \tag{3.5.38}$$

Theorem 3.5.39 *(LPIT2) The assumption 3.5.5 (FBAss2) can be satisfied for each compact domain $\Omega \subset I\!R^d$ satisfying an interior cone condition 3.5.32 (DefICC) that defines positive constants $\epsilon_c,\ K_c$. If γ_d is defined as in 3.5.34 (MNL), the constants C_2 and C_3 can be bounded by (3.5.38, C2def) and (3.5.36, C3def), respectively, while the polynomial order ℓ and the fill distance h must satisfy (3.5.37, hrestr).*

3.5.8 Error Bounds in Terms of Fill Distance

(*hrhodef*) We can now assemble the previous results into bounds of the form (3.4.15, *FBound*) for the power function from optimal recovery. Together with (3.4.4, *EqgSg1*) from page 64 this yields error bounds for the reconstruction of functions g from native spaces $\mathcal{G}$. Depending on the situation, we get quite explicit bounds for the power function in cases of small space dimensions and polynomial orders, while for fixed orders and arbitrary space dimension we use Theorem 3.5.33 (*LPIT)* to carry the order of the local bounds on the power functions over to the errors of optimal recoveries, the constants being mysterious. We list the orders (without the factors) of our L_∞ bounds on the power function in Table 6 (*TCPDEB)*, but delay the cases with exponential convergence somewhat. The additional data (parameters, domains, smoothness, dimension, order) should be looked up from tables 1 (*TCPDFct)* and 2 (*TPDFct)* on page 18. Note that the actual approximation orders of optimal recoveries may be better than the squares of these bounds.

$\phi(r)$	L_∞ Bound of Power Function
r^β	$h^{\beta/2}$
$r^\beta \log r$	$h^{\beta/2}$
$(r^2+\gamma^2)^{\beta/2}$	$\exp(-c/h),\ c>0$
$e^{-\beta r^2}$	$\exp(-c/h^2),\ c>0$
$r^\nu K_\nu(r)$	h^ν
$(1-r)_+^2(2+r)$	$h^{1/2}$
$(1-r)_+^4(1+4r)$	$h^{3/2}$

Table 6: L_∞ Bounds of Power Function Based on Lagrange Data (*TCPDEB*)

Unfortunately, the factor $|g-S(g)|_\Phi$ in the actual error bound (3.4.4, *EqgSg1*) still is somewhat mysterious, if we start with a conditionally positive definite function Φ and construct the corresponding native space. If, on the other hand, we have started with $\mathcal{G}$, we are done. But note that these bounds can be improved, if g satisfies additional conditions. These improvements cannot come from better bounds on the power function, because we shall see that our techniques often provide optimal orders with respect to h. They rely on a deeper analysis of the term $|g-S(g)|_\Phi$, and this analysis will be done in 3.6 (*SecEBStage2)* and 3.7 (*SecEBStage3)*.

We now discuss the cases of multiquadrics and Gaussians, where we can push the polynomial order ℓ up to infinity. The overall bound for $P^2(x)$ is given by (3.5.10, *FundBound*), and we have to insert (3.5.38, *C2def*), (3.5.36, *C3def*) and the replacement of h by $2h\sqrt{d}$ from section 3.5.2 (*SecAERC)*. The values of ρ and C_1 depend on the special case chosen.

Let us first look at multiquadrics. The bound on $P^2(x)$ then is

$$P^2(x) \;\leq\; 2^{|\beta/2|+1}c^\beta \left(1+e^{2d\ell\gamma_d}\right)^2 \left(4(K_c+1)\ell\sqrt{d}\gamma_d\right)^{2n} \left(\tfrac{8h^2d}{c^2}\right)^n$$

under the restrictions $8h^2d < c^2/2$ and (3.5.37, *hrestr*). We now treat everything as fixed except h and $\ell = 2n-1$. This turns the bound into something of the form

$$C_4\left(C_5nh\right)^{2n},$$

and we shall pick $n = (\ell+1)/2$ as a function of h as large as possible, but such that the constraints

$$C_5nh \leq \gamma < 1,\; 4\ell h\gamma_d \leq \epsilon_c$$

are satisfied. This works with $n = ch^{-1}/2$ and some positive constant c. Then the bound becomes

$$C_4\gamma^{c/h} = C_4\exp(-|\log\gamma|c/h)$$

and proves exponential behaviour for $h \to \infty$.

The Gaussian case is quite similar and can easily be reduced to a bound like

$$C_4\frac{\left(C_5nh\right)^{2n}}{n!},$$

which allows the same treatment. But now we can use the additional $n!$ in the denominator to speed up the convergence. We first insert Stirling's formula (3.5.27, *Stirling*) into the denominator to cancel an n^n factor in the numerator, introducing some change in the constants C_4 and C_5. This yields

$$C_4\left(C_5\sqrt{n}h\right)^{2n},$$

and we now pick $n = ch^{-2}/2$ to get

$$C_5\sqrt{n}h \leq \gamma < 1.$$

The second restriction, induced by the interior cone condition, cannot be satisfied in this case, and this is why our final result will only hold in the interior or near smooth boundary parts of Ω. We get the bound

$$C_4\gamma^{c/h^2} = C_4\exp(-|\log\gamma|c/h^2)$$

with "Gaussian" exponential behaviour for $h \to \infty$.

Theorem 3.5.40 *(GMCEBT) The power functions of Lagrange interpolation by multiquadrics and Gaussians have L_∞ bounds of the form $\exp(-c/h)$ with $c > 0$ for compact domains $\Omega \subset I\!R^d$ satisfying an interior cone condition. The bound for the Gaussian can be improved to $\exp(-c/h^2)$ in the interior or near smooth parts of the boundary of Ω.*

3.5.9 Remarks

The proof of bounds on the power function via polynomial approximation goes back to Duchon [6](*duchon:76-1*) for thin-plate splines and was successfully generalized by Madych and Nelson [15](*madych-nelson:88-1*), [18](*madych-nelson:92-1*). The special cases of lines and triangles were done for thin-plate splines by Powell [27](*powell:93-1*).

3.6 Doubling the Approximation Order

(*SecEBStage2*) Here we show how the error bounds of the form (3.4.4, *EqgSg1*) can be improved by adding some assumptions on the function g that is reconstructed. A third enhancement, based on a localization argument, will follow in 3.7 (*SecEBStage3).*

We work in the setting of section 3.3.5 (*SecSN)* and define the bilinear form (*Norm2*)

$$(f,g)_2 := \int_\Omega \delta_{x,\mathcal{P}} f \cdot \delta_{x,\mathcal{P}} g dx \tag{3.6.1}$$

for all $f, g \in \mathcal{G}$. Because of

$$(\delta_{x,\mathcal{P}} f)^2 = (\delta^y_{x,\mathcal{P}} \Phi(y,\cdot), f)^2_\Phi \le \|\delta_{x,\mathcal{P}}\|^2_\Phi \|f\|^2_\Phi$$

and since $\Psi_\mathcal{P}(x,y) = (\delta_{x,\mathcal{P}}, \delta_{y,\mathcal{P}})_\Phi$ is the normalization of Φ with respect tp $\mathcal{P}$, we can use its continuity and get that the bilinear form (3.6.1, *Norm2*) is well-defined and continuous with respect to the bilinear form in $\mathcal{G}$:

$$|(f,g)_2| \le \|f\|_2 \|g_2 \le \|\Psi_\mathcal{P}(x,x)\|_{L_2(\Omega)} \|f\|_\Phi \|g\|_2 \le \|\Psi_\mathcal{P}(x,x)\|^2_{L_2(\Omega)} \|f\|_\Phi \|g\|_\Phi.$$

We apply this to (3.4.4, *EqgSg1*) and get

$$\|g - S(g)\|_2 \le \|P_\Lambda(\cdot)\|_{L_2(\Omega)} \|g - S(g)\|_\Phi$$

for the optimal power function P_Λ and the optimal recovery $S(g)$ of $g \in \mathcal{G}$.

For any $g \in \mathcal{G}$ we can consider the continuous linear functional $f \mapsto (f,g)_2$ on the Hilbert space $\mathcal{F}$. Then there is a function $Cg \in \mathcal{G}$ such that

(*Cmapdef*)

$$(f,g)_2 = (f,Cg)_\Phi \text{ for all } f \in \mathcal{G}. \tag{3.6.2}$$

We now define the subspace

$$\mathcal{H} := \mathcal{P} + C(\mathcal{F}) \subseteq \mathcal{G}$$

of $\mathcal{G}$ and consider optimal recovery of functions $g = p + C(f_g) \in \mathcal{H}$ by $g^* = S(g)$. The orthogonality (3.1.34, *EqOrtho*) then implies

$$\begin{aligned}
\|g - S(g)\|_\Phi^2 &= (g - S(g), g - S(g))_\Phi \\
&= (g, g - S(g))_\Phi \\
&= (Cf_g, g - S(g))_\Phi \\
&= (f_g, g - S(g))_2 \\
&\leq \|f_g\|_2 \|g - S(g)\|_2 \\
&\leq \|f_g\|_2 \|P_\Lambda(\cdot)\|_{L_2(\Omega)} \|g - S(g)\|_\Phi
\end{aligned}$$

and this allows to bound $\|g - S(g)\|_\Phi$ nicely by

$$\|g - S(g)\|_\Phi \leq \|P_\Lambda\|_{L_2(\Omega)} \|f_g\|_2$$

for all $g = p + Cf_g \in \mathcal{H}$. If we combine this with (3.4.4, *EqgSg1*), we get

Theorem 3.6.3 *(EBStage2T) For optimal reconstruction of functions $g \in \mathcal{H} \subset \mathcal{G}$ with $g = Cf_g$ in the sense of (3.6.2, Cmapdef) by optimal recovery functions $S(g)$ we have the improved error bound*

$$|(g - S(g))(x)| \leq P(x) \|P\|_{L_2(\Omega)} \|f_g\|_2$$

for all $x \in \Omega$.

3.6.1 Remarks

The results of 3.6 (*SecEBStage2)* are from [9](*RSImpEB*) and derived from the arguments used in classical spline theory [1](*ahlberg-et-al:86-1*) to improve the approximation order via the "second integral relation".

3.7 Improvement by Localization

(*SecEBStage3*) Here we use a localization argument dating back to Duchon [6](*duchon:76-1*) to get some additional powers of h for error bounds of optimal recovery via Lagrange interpolation. We delay the formulation of these results until the current discussion between M. Buhmann, W. Light, and R. S. about this topic has settled.

4 Advanced Results on $\mathbb{R}^d$

(*SecARRd*) Here we introduce Fourier transforms in $\mathbb{R}^d$, and derive a series of results that require related techniques. These include bounds on the stability and error bounds for the multilevel method.

4.1 Fourier Transforms on $\mathbb{R}^d$

4.1.1 Fourier Transforms of Tempered Test Functions

There are two major possibilities to pick a space $\mathcal{S}$ of test functions on $\mathbb{R}^d$ to start with, and we take the **tempered test functions** that are verbally defined as real-valued functions on $\mathbb{R}^d$ whose partial derivatives exist for all orders and decay faster than any polynomial towards infinity.

Definition 4.1.1 *(DefFT) For a test function $u \in \mathcal{S}$, the* **Fourier transform** *is*

$$\widehat{u}(\omega) := (2\pi)^{-d/2} \int_{\mathbb{R}^d} u(x) e^{-ix\cdot\omega} dx,$$

where ω varies in $\mathbb{R}^d$ and $x \cdot \omega$ is shorthand for the scalar product $x^T\omega = \omega^T x$ to avoid the ^{T} symbol in the exponent. Since the definition even works for general $u \in L_1(\mathbb{R}^d)$, it is well-defined on $\mathcal{S}$ and clearly linear. Note that we use the **symmetric** form of the transform and do not introduce a factor 2π in the exponent of the exponential. This sometimes makes comparisons to other presentations somewhat difficult.

To get used to calculations of Fourier transforms, let us start with the **Gaussian** $u_\gamma(x) = \exp(-\gamma\|x\|_2^2)$ for $\gamma > 0$, which clearly is in the space of test functions, since all derivatives are polynomials multiplied with the Gaussian itself. As a byproduct we shall get that the Gaussian is positive definite on $\mathbb{R}^d$. Fortunately, the Gaussian can be written as a d-th power of the entire analytic function $\exp(-\gamma z^2)$, and we can thus work on $\mathbb{C}^d$ instead of $\mathbb{R}^d$. We simply use substitution in

$$\begin{aligned}
\widehat{u_\gamma}(i\omega) &= (2\pi)^{-d/2} \textstyle\int_{\mathbb{R}^d} e^{-\gamma\|x\|_2^2} e^{x\cdot\omega} dx \\
&= (2\pi)^{-d/2} e^{\|\omega\|_2^2/4\gamma} \textstyle\int_{\mathbb{R}^d} e^{-\|\sqrt{\gamma}x - \omega/2\sqrt{\gamma}\|_2^2} dx \\
&= (2\pi\gamma)^{-d/2} e^{\|\omega\|_2^2/4\gamma} \textstyle\int_{\mathbb{R}^d} e^{-\|y\|_2^2} dy
\end{aligned}$$

and are done up to the evaluation of the dimension-dependent constant

$$\int_{\mathbb{R}^d} e^{-\|y\|_2^2} dy =: c^d$$

which is a d-th power, because the integrand factorizes nicely. We calculate c^2 by using polar coordinates and get

$$\begin{aligned} c^2 &= \int_{\mathbb{R}^2} e^{-\|y\|_2^2} dy \\ &= \int_0^{2\pi} \int_0^\infty e^{-r^2} r \, dr \, d\varphi \\ &= 2\pi \int_0^\infty e^{-r^2} r \, dr \\ &= -\pi \int_0^\infty (-2r) e^{-r^2} \, dr \\ &= \pi. \end{aligned}$$

This proves the first assertion of

Theorem 4.1.2 *(GaussPD) The Gaussian*

$$u_\gamma(x) = \exp(-\gamma \|x\|_2^2)$$

has Fourier transform

(GFT)

$$\widehat{u_\gamma}(\omega) = (2\gamma)^{-d/2} e^{-\|\omega\|_2^2/4\gamma} \tag{4.1.3}$$

and is unconditionally positive definite on $\mathbb{R}^d$.

Proof: Let us first invert the Fourier transform by setting $\beta := 1/4\gamma$ in (4.1.3, *GFT*):

$$\begin{aligned} \exp(-\beta \|\omega\|_2^2) &= (4\pi\beta)^{-d/2} \int_{\mathbb{R}^d} e^{-\|x\|_2^2/4\beta} e^{-ix\cdot\omega} dx \\ &= (2\pi)^{-d/2} \int_{\mathbb{R}^d} (2\beta)^{-d/2} e^{-\|x\|_2^2/4\beta} e^{+ix\cdot\omega} dx. \end{aligned}$$

Then take any set $X = \{x_1, \ldots, x_M\} \subset \mathbb{R}^d$ of M distinct points and any vector $\alpha \in \mathbb{R}^M$ to form

$$\begin{aligned} \alpha^T A_{X,u_\gamma} \alpha &= \sum_{j,k=1}^M \alpha_j \alpha_k \exp(-\gamma \|x_j - x_k\|_2^2) \\ &= \sum_{j,k=1}^M \alpha_j \alpha_k (4\pi\gamma)^{-d/2} \int_{\mathbb{R}^d} e^{-\|x\|_2^2/4\gamma} e^{-ix\cdot(x_j - x_k)} dx \\ &= (4\pi\gamma)^{-d/2} \int_{\mathbb{R}^d} e^{-\|x\|_2^2/4\gamma} \sum_{j,k=1}^M \alpha_j \alpha_k e^{-ix\cdot(x_j - x_k)} dx \\ &= (4\pi\gamma)^{-d/2} \int_{\mathbb{R}^d} e^{-\|x\|_2^2/4\gamma} \left| \sum_{j=1}^M \alpha_j e^{-ix\cdot x_j} \right|^2 dx \geq 0. \end{aligned}$$

This proves positive semidefiniteness of the Gaussian. To prove definiteness, we can assume

$$f(x) := \sum_{j=1}^{M} \alpha_j e^{-ix\cdot x_j} = 0$$

for all $x \in \mathbb{R}^d$ and have to prove that all coefficients α_j vanish. Taking derivatives at zero, we get

$$0 = D^\beta f(0) = \sum_{j=1}^{M} \alpha_j (-ix_j)^\beta,$$

and this is a homogeneous system for the coefficients α_j whose coefficient matrix is a generalized Vandermonde matrix, possibly transposed and with scalar multiples for rows or columns. This proves the assertion in one dimension, where the matrix corresponds to the classical Vandermonde matrix. The multivariate case reduces to the univariate case by picking a nonzero vector $y \in \mathbb{R}^d$ that is not orthogonal to any of the finitely many differences $x_j - x_k$ for $j \neq k$. Then the real values $y \cdot x_j$ are all distinct for $j = 1, \ldots, M$ and one can consider the univariate function

$$g(t) := f(ty) = \sum_{j=1}^{M} \alpha_j e^{-ity\cdot x_j} = 0$$

which does the job in one dimension. □

Note that the Gaussian is mapped to itself by the Fourier transform, if we pick $\gamma = 1/2$. We shall use the Gaussian's Fourier transform in the proof of the fundamental **Fourier Inversion Theorem**:

Theorem 4.1.4 *(FTTS) The Fourier transform is bijective on $\mathcal{S}$, and its inverse is the transform*

$$\check{u}(x) := (2\pi)^{-d/2} \int_{\mathbb{R}^d} u(\omega) e^{ix\cdot\omega} d\omega.$$

Proof: The multivariate derivative D^α of $\widehat{u}$ can be taken under the integral sign, because u is in $\mathcal{S}$. Then

$$(D^\alpha \widehat{u})(\omega) = (2\pi)^{-d/2} \int_{\mathbb{R}^d} u(x)(-ix)^\alpha e^{-ix\cdot\omega} dx,$$

and we multiply this by ω^β and use integration by parts

$$\begin{aligned}
\omega^\beta (D^\alpha \widehat{u})(\omega) &= (2\pi)^{-d/2} \textstyle\int_{\mathbb{R}^d} u(x)(-ix)^\alpha (i)^\beta (-i\omega)^\beta e^{-ix\cdot\omega} dx \\
&= (2\pi)^{-d/2} \textstyle\int_{\mathbb{R}^d} u(x)(-ix)^\alpha (i)^\beta \frac{d^\beta}{dx^\beta} e^{-ix\cdot\omega} dx \\
&= (2\pi)^{-d/2} (-1)^{|\alpha|+|\beta|} i^{\alpha+\beta} \textstyle\int_{\mathbb{R}^d} e^{-ix\cdot\omega} \frac{d^\beta}{dx^\beta} (u(x)x^\alpha) dx
\end{aligned}$$

to prove that $\widehat{u}$ lies in $\mathcal{S}$, because all derivatives decay faster than any polynomial towards infinity. The second assertion follows from the Fourier inversion formula

$$u(x) := (2\pi)^{-d/2} \int_{\mathbb{R}^d} \widehat{u}(\omega) e^{ix\cdot\omega} d\omega$$

that we now prove for all $u \in \mathcal{S}$. This does not work directly if we naively put the definition of $\widehat{u}$ into the right-hand-side, because the resulting multiple integral does not satisfy the assumptions of Fubini's theorem. We have to do a regularization of the integral, and since this is a standard trick, we write it out in some detail:

$$\begin{aligned}
(2\pi)^{-d/2} \textstyle\int_{\mathbb{R}^d} \widehat{u}(\omega) e^{ix\cdot\omega} d\omega &= (2\pi)^{-d} \textstyle\int_{\mathbb{R}^d} \int_{\mathbb{R}^d} u(y) e^{i(x-y)\cdot\omega} dy d\omega \\
&= \lim_{\epsilon \searrow 0} (2\pi)^{-d} \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} u(y) e^{i(x-y)\cdot\omega - \epsilon\|\omega\|_2^2} dy d\omega \\
&= \lim_{\epsilon \searrow 0} (2\pi)^{-d} \int_{\mathbb{R}^d} \left(\int_{\mathbb{R}^d} e^{i(x-y)\cdot\omega - \epsilon\|\omega\|_2^2} d\omega \right) u(y) dy \\
&= \lim_{\epsilon \searrow 0} \int_{\mathbb{R}^d} \varphi(\epsilon, x-y) u(y) dy
\end{aligned}$$

with

(*deltaschar*)

$$\varphi(\epsilon, z) := (2\pi)^{-d} \int_{\mathbb{R}^d} e^{iz\cdot\omega - \epsilon\|\omega\|_2^2} d\omega. \tag{4.1.5}$$

The proof is completed by application of the following result that is useful in many contexts: □

Lemma 4.1.6 *(LemRepro) The family of functions $\varphi(\epsilon, z)$ of (4.1.5, deltaschar) approximates the point evaluation functional in the sense*

(Repro)

$$u(x) = \lim_{\epsilon \searrow 0} \int_{\mathbb{R}^d} \varphi(\epsilon, x-y) u(y) dy \tag{4.1.7}$$

for all functions u that are in $L_1(\mathbb{R}^d)$ and continuous around x.

Proof: We first remark that φ is a disguised form of the inverse Fourier transform equation of the Gaussian. Thus we get

(*deltarep*)

$$\varphi(\epsilon, x) = (4\pi\epsilon)^{-d/2} e^{-\|x\|_2^2/4\epsilon} \tag{4.1.8}$$

and

$$\int_{\mathbb{R}^d} \varphi(\epsilon, x) dx = (4\pi\epsilon)^{-d/2} \int_{\mathbb{R}^d} e^{-\|x\|_2^2/4\epsilon} dx = 1.$$

To prove (4.1.7, *Repro*), we start with some given $\delta > 0$ and first find some ball $B_\rho(x)$ of radius $\rho(\delta)$ around x such that $|u(x) - u(y)| \leq \delta/2$ holds uniformly for all $y \in B_\rho(x)$. Then we split the integral in

$$\begin{array}{rcl} |u(x) - \int_{\mathbb{R}^d} \varphi(\epsilon, x-y)u(y)dy| & = & |\int_{\mathbb{R}^d} \varphi(\epsilon, x-y)(u(x)-u(y))dy| \\ & \leq & \int_{\|y-x\|_2 \leq \rho} \varphi(\epsilon, x-y)|u(x)-u(y)|dy \\ & & + \int_{\|y-x\| > \rho} \varphi(\epsilon, x-y)|u(x)-u(y)|dy \\ & \leq & \delta/2 + (4\pi\epsilon)^{-d/2} e^{-\rho^2/4\epsilon} 2\|u\|_1 \\ & \leq & \delta \end{array}$$

for all sufficiently small ϵ. □

Due to the Fourier inversion formula, we now know that the Fourier transform is bijective on $\mathcal{S}$.

We now relate the Fourier transform to the L_2 inner product, but we have to use the latter over $\mathbb{C}$ to account for the possibly complex values of the Fourier transform. Furthermore, we have good reasons to define the inner product as

(*Ltwodef*)

$$(f,g)_{L_2(\mathbb{R}^d)} := (2\pi)^{-d/2} \int_{\mathbb{R}^d} f(x)\overline{g(x)}dx \tag{4.1.9}$$

with a factor that simplifies some of the subsequent formulae.

Fubini's theorem easily proves the identity

$$(v, \widehat{u})_{L_2(\mathbb{R}^d)} = (2\pi)^{-d} \int_{\mathbb{R}^d} v(x) \int_{\mathbb{R}^d} \overline{u(y)} e^{+ix\cdot y} dydx = (\check{v}, u)_{L_2(\mathbb{R}^d)}$$

for all test functions $u, v \in \mathcal{S}$. Setting $v = \widehat{w}$ we get Parceval's equation

(*ParsRd*)

$$(\widehat{w}, \widehat{u})_{L_2(\mathbb{R}^d)} = (w, u)_{L_2(\mathbb{R}^d)} \tag{4.1.10}$$

for the Fourier transform on $\mathcal{S}$, proving that the Fourier transform is isometric on $\mathcal{S}$ as a subspace of $L_2(\mathbb{R}^d)$.

4.1.2 Fourier Transforms of Functionals

With Parceval's equation in mind, let us look at the linear functional

$$\lambda_u(v) := (u, v)_{L_2(\mathbb{R}^d)}$$

on $\mathcal{S}$. We see that

$$\lambda_{\widehat{u}}(v) = (\widehat{u}, v)_{L_2(\mathbb{R}^d)} = (u, \check{v})_{L_2(\mathbb{R}^d)} = \lambda_u(\check{v})$$

holds. A proper definition of the Fourier transform for functionals λ_u should be in line with the functions u that represent them, and thus we should define

$$\widehat{\lambda_u} := \lambda_{\widehat{u}}$$

or in more generality

$$\widehat{\lambda}(v) := \lambda(\check{v})$$

for all $v \in \mathcal{S}$. Since the space $\mathcal{S}$ of test functions is quite small, its dual, the space of linear functionals on $\mathcal{S}$, will be quite large.

Definition 4.1.11 *The Fourier transform of a linear functional λ on $\mathcal{S}$ is the linear functional $\widehat{\lambda}$ on $\mathcal{S}$ defined by*

$$\widehat{\lambda}(v) := \lambda(\check{v})$$

for all $v \in \mathcal{S}$.

If we can represent the functional $\widehat{\lambda}$ as λ_v, we write $v = \widehat{\lambda}$ as a shorthand notation, but keep the original meaning in mind. Let us look at some examples.

Example 4.1.12 *(ExDelta) The functional $\delta_x(v) := v(x)$ has the form*

$$\delta_x(v) = v(x) = (2\pi)^{-d/2} \int_{\mathbb{R}^d} \widehat{v}(\omega) e^{+ix\cdot\omega} d\omega,$$

and its Fourier transform is of the form λ_u with

$$u(\omega) = \widehat{\delta_x}(\omega) = e^{-x\cdot\omega}.$$

Here, the normalization of the L_2 inner product (4.1.9, Ltwodef) pays off. Note that the Fourier transform is not a test function, but rather an at most polynomially growing function from $\mathcal{K}$ and in particular a bounded function. The functional $\delta := \delta_0$ has the Fourier transform 1.

Example 4.1.13 *(Exlxma) A very important class of functionals for our purposes consists of the space $\mathcal{P}_\Omega^\perp$ of functionals of the form (3.3.1, Deflxma) that vanish on $\mathbb{P}_m^d$. Their action on a test function v is*

$$\begin{aligned} \lambda_{X,M\alpha}(v) &= \sum_{j=1}^{M} \alpha_j v(x_j) \\ &= (2\pi)^{-d/2} \int_{\mathbb{R}^d} \widehat{v}(\omega) \sum_{j=1}^{M} \alpha_j e^{ix_j\cdot\omega} d\omega \\ &= \widehat{\lambda}_{X,M\alpha}(\widehat{v}) \end{aligned}$$

such that the Fourier transform of the functional $\lambda_{X,M\alpha}$ is the functional generated by the bounded function

$$\widehat{\lambda}_{X,M,\alpha}(\omega) = \sum_{j=1}^{M} \alpha_j e^{-ix_j\cdot\omega}.$$

If we expand the exponential into its power series, we see that

$$\begin{aligned}\widehat{\lambda}_{X,M,\alpha}(\omega) &= \sum_{k=0}^{\infty}\sum_{j=1}^{M} \alpha_j(-ix_j\cdot\omega)^k/k! \\ &= \sum_{k=m}^{\infty}\sum_{j=1}^{M} \alpha_j(-ix_j\cdot\omega)^k/k!\end{aligned}$$

since the functional vanishes on $\mathbb{P}_m^d$. Thus $\widehat{\lambda}_{X,M,\alpha}(\omega)$ has a zero of order at least m in zero. If the functional $\lambda_{X,M\alpha}$ itself were representable by a function u, the function u should be orthogonal to all polynomials from $\mathbb{P}_m^d$. We shall use both of these facts later.

Example 4.1.14 *(ExFTPol) The monomials x^α are in the space $\mathcal{K}$, and thus they should at least have generalized Fourier transforms in the sense of functionals. This can easily be verified via*

$$\begin{aligned}\left(-i\tfrac{d}{dx}\right)^\alpha v(x) &= \left(-i\tfrac{d}{dx}\right)^\alpha (2\pi)^{-d/2}\int_{\mathbb{R}^d}\widehat{v}(\omega)e^{+ix\cdot\omega}d\omega \\ &= (2\pi)^{-d/2}\int_{\mathbb{R}^d}\widehat{v}(\omega)(-i\cdot i\omega)^\alpha e^{+ix\cdot\omega}d\omega \\ &= (2\pi)^{-d/2}\int_{\mathbb{R}^d}\widehat{v}(\omega)\omega^\alpha e^{+ix\cdot\omega}d\omega\end{aligned}$$

and the associated functional is

$$v \mapsto \left(-i\frac{d}{dx}\right)^\alpha v(x)$$

at $x = 0$.

4.1.3 Fourier Transform in $L_2(\mathbb{R}^d)$

The test functions from $\mathcal{S}$ are dense in $L_2(\mathbb{R}^d)$ (see Lemma 10.5.3 (*FTD*) for details), and thus we have

Theorem 4.1.15 *(FLtwoT) The Fourier transform has an L_2-isometric extension from the space $\mathcal{S}$ of tempered test functions to $L_2(\mathbb{R}^d)$. The same holds for the inverse Fourier transform, and both extensions are inverses of each other in $L_2(\mathbb{R}^d)$. Furthermore, Parceval's equation (4.1.10, ParsRd) holds in $L_2(\mathbb{R}^d)$.* □

Note that this result does not allow to use the Fourier transform formula (or its inverse) in the natural pointwise form. For any $f \in L_2(\mathbb{R}^d)$ one first has to provide a sequence of test functions $v_n \in \mathcal{S}$ that converges to f in the L_2 norm for $n \to \infty$, and then, by continuity, the image $\widehat{f}$ of the Fourier transform is uniquely defined almost everywhere by

$$\lim_{n\to\infty} \|\widehat{f} - \widehat{v_n}\|_{L_2(\mathbb{R}^d)} = 0.$$

This can be done via Friedrich's mollifiers as defined in (10.5.2, *Friedmoll*), replacing the Gaussian in the representation (4.1.8, *deltarep*) by a compactly supported infinitely differentiable function.

A more useful characterization of $\widehat{f}$ is the variational equation

$$(\widehat{f}, v)_{L_2(\mathbb{R}^d)} = (f, \check{v})_{L_2(\mathbb{R}^d)}$$

for all test functions $v \in \mathcal{S}$, or, by continuity, all functions $v \in L_2(\mathbb{R}^d)$.

4.2 Native Spaces Characterized by Transforms

(*SecCNST*) In section 3.2.4 (*SecIP*) we have seen that on $\mathbb{R}^d$ we can restrict ourselves to cases where the recovery problem is translation-invariant or even invariant under Euclidean rigid-body transformations. In the first case, the conditionally positive definite functions $\Phi(x, y)$ take the form $\Phi(x-y)$, while in the second they are radial: $\Phi(x, y) = \phi(\|x-y\|_2)$. We start with the more general case, but we restrict ourselves to **un**conditionally positive definite functions first.

So let us now consider a function $\Phi : \mathbb{R}^d \to \mathbb{R}$ with $\Phi(-\cdot) = \Phi(\cdot)$ such that $\Psi(x, y) := \Phi(x - y)$ is an unconditionally positive definite function on $\mathbb{R}^d$. Having the Gaussian in mind as a prominent example, we assume Φ to have a continuous nonnegative real-valued and integrable Fourier transform $\widehat{\Phi}$ on $\mathbb{R}^d$ such that the Fourier inversion formula holds. We now want to construct the native space $\mathcal{G}$ by the techniques of section 3.3 (*SecNS*). Clearly, the representation (3.3.2, *DefBil*) of the bilinear form $(\cdot, \cdot)_\Phi$ can now be rewritten as

(*DefBil2*)

$$(\lambda_{X,M,\alpha}, \lambda_{Y,N,\beta})_\Phi := (2\pi)^{-d/2} \int_{\mathbb{R}^d} \widehat{\Phi}(\omega) \sum_{j=1}^{M} \sum_{k=1}^{N} \alpha_j \beta_k e^{i(x_j - y_k)\cdot\omega} \Phi d\omega. \quad (4.2.1)$$

The functions $F(\lambda_{X,M,\alpha})$ have Fourier transforms

$$\begin{aligned} \widehat{F}(\lambda_{X,M,\alpha})(\omega) &= \widehat{\Phi}(\omega)\sum_{j=1}^{M}\alpha_j e^{-x_j\cdot\omega} \\ &= \widehat{\Phi}(\omega)\widehat{\lambda}_{X,M,\alpha} \end{aligned}$$

if we define Fourier transforms of functionals appropriately.

4.3 Condition Numbers

(*SecCNTrans*)

4.4 Error Bounds

(*SecEBTrans*)

Example 4.4.1 *(EBSob) Sobolev radial bsis functions*

5 Special Theory

(*SecST*) Here we introduce general transforms and generalize the results that we had on $I\!R^d$. We start with generalizing the notion of a transform in order to cover some other cases we consider in some detail later:

1. Fourier series on $[0, 2\pi]^d$,
2. General expansions in orthogonal series,
3. Harmonic analysis on locally compact topological groups.

It will turn out that certain results can be formulated for general transforms, while others take advantage of the special structure of the underlying space.

5.1 Results for General Transforms

(*SecT*) This section covers the necessary results about general transforms. The applications except for R^d will follow later. We start from the general setting and add the specific details later.

5.1.1 General Transforms

(*SecGTR*) Here we formulate the general notions that apply to all kinds of transforms that we consider later. The setting is general enough to allow **generalized** transforms in addition to the classical ones. This turns out to be absolutely necessary even in the simple case of Fourier transforms on $\mathbb{R}^d$. For this reason we do not rely on other sources on transforms.

Assume that our basic space $\mathcal{G}$ with positive definite bilinear form $(\cdot,\cdot)_\Phi$ and nullspace $\mathcal{P}$ is a space of real-valued functions on some domain Ω. Forget about $\Phi, \mathcal{G}$, and $\mathcal{P}$ for a moment, and concentrate on Ω.

Assumption 5.1.1 *(TAss1) For a specific space $\mathcal{S}$ of real-valued* **test functions** *on some domain Ω there is a linear and injective* **transform mapping**

$$g \mapsto \widehat{g} \ : \ \mathcal{S} \to L_{2,\sigma}(D)$$

whose values are complex-valued functions on some domain D that carries a measure σ such that the space

$$L_{2,\sigma}(D) := \left\{ f \ : \ D \to \mathbb{C} \ : \ \int_D |f|^2 d\sigma \ < \infty \right\}$$

is well-defined and a Hilbert space over $\mathbb{C}$ with inner product

$$(u,v)_{L_{2,\sigma}(D)} := \int_D u\overline{v} d\sigma \text{ for all } u, v \in L_{2,\sigma}(D).$$

In particular, the measure σ can be Lebesgue measure on $D = \mathbb{R}^d$ for the classical multivariate Fourier transform, or Haar measure on a locally compact topological group, or plain summation for series transforms, e.g.: $D = \mathbb{Z}^d$ for Fourier series on $\Omega = [0, 2\pi]^d$. Note that the transform domain D and its measure σ are independent of the functions Φ that we are going to consider, but they will crucially depend on Ω. We shall often write $(\cdot,\cdot)_2$ as shorthand for the above inner product, and we use the phrase **almost everywhere** to stand for "on D except for a set of σ-measure zero". Assumption 5.1.1 *(TAss1)* is usually satisfied by proper definition of D, σ, and the transform mapping. Injectivity of the latter is often proved by an inverse transform.

Assumption 5.1.2 *(TAss2) The space $L_{2,\sigma}(D)$ contains the image of the space $\mathcal{S}$ of test functions under the transform mapping as a dense subspace and coincides with its closure under the inner product $(\cdot,\cdot)_{L_{2,\sigma}(D)}$.*

This makes sure that the test function space $\mathcal{S}$ is rich enough to generate all of $L_{2,\sigma}(D)$ by continuity arguments acting on transforms.

Assumption 5.1.3 *(TAss3) There is a 1-1 correspondence between L_2 spaces on Ω and D in the sense that there is a measure ω on Ω such that the spaces $L_{2,\sigma}(D)$ and $L_{2,\omega}(\Omega)$ are isometrically isomorphic under the transform mapping:*

(Pars)

$$(f,g)_{L_{2,\omega}\Omega} := \int_\Omega f\overline{g}d\omega = (\widehat{f},\overline{\widehat{g}})_{L_{2,\sigma}(D)}. \tag{5.1.4}$$

Identities like (5.1.4, *Pars*) are usually called **Parseval's equation**. Of course, one could use the structure on $L_{2,\sigma}(D)$ to define an inner product for functions on Ω by using (5.1.4, *Pars*) without the representation via integrals as a definition. Thus the actual meaning of Assumption 5.1.3 (*TAss3)* is that this abstract inner product can be respresented as a standard L_2 inner product.

5.1.2 Spaces Induced by Basis Functions

We restrict ourselves to basis functions Φ that satisfy

Assumption 5.1.5 *(PFTAss1) The conditionally positive definite function $\Phi : \Omega\times\Omega \to I\!R$ has an associated real-valued nonnegative function $\widehat{\Phi}$ which is defined and positive almost everywhere on the transform domain D.*

For reasons to become apparent later, we do not require $\widehat{\Phi}$ to be the image of Φ under the transform mapping, since we shall encounter cases where Φ is not in the domain of the transform. One should rather consider $(\widehat{\Phi})^{-1}$ as a weight function on D. But there will also be cases where actually $\widehat{\Phi}$ is the transform of Φ, thus the notation. The relation between Φ and $\widehat{\Phi}$ will be clarified after introducing some additional notation.

We use $(\widehat{\Phi})^{-1}$ as a weight function to define the operator

$$L_\Phi \ : \ g \mapsto \frac{\widehat{g}}{\sqrt{\widehat{\Phi}}}.$$

To turn it into a continuous map with image in $L_{2,\sigma}(D)$, we restrict its domain to the subspace

$$\mathcal{S}_\Phi := \{u \in \mathcal{S} \ : \ L_\Phi u \in L_{2,\sigma}(D)\}$$

of the space $\mathcal{S}$ of test functions on D. We now can define an inner product

(PTIP)

$$(f,g)_\Phi = \int_D \widehat{f}(\widehat{\Phi})^{-1}\overline{\widehat{g}}d\sigma \tag{5.1.6}$$

on all $f, g \in \mathcal{S}_\Phi$.

We are now ready to link $\widehat{\Phi}$ to Φ and its native space $\mathcal{G} = \mathcal{P} + \mathcal{F}$ by the requirement

Assumption 5.1.7 *(PTAss2) The closure of $\mathcal{S}_\Phi$ under the inner product (5.1.6, PTIP) coincides with the Hilbert space $\mathcal{F}$.*

Then the mapping L_Φ can be identified with its continuous extension to all of $\mathcal{F}$, and it can be further extended to $\mathcal{G}$ by defining it as being zero on $\mathcal{P}$. The image of $\mathcal{F}$ under L_Φ is a closed Hilbert subspace of $L_{2,\sigma}(D)$, and we shall require some additional work in special cases to prove

Assumption 5.1.8 *(PTAss3) The mapping*

(LSurj)

$$L_\Phi \ : \ \mathcal{G} \to L_{2,\sigma}(D) \tag{5.1.9}$$

as the canonical extension of

$$L_\Phi(g) := \frac{\widehat{g}}{\sqrt{\widehat{\Phi}}}$$

for $g \in \mathcal{S}_\Phi$ is surjective.

The extension allows to define a **generalized transform** on the space $\mathcal{G}$ via

$$\widehat{g} := \sqrt{\widehat{\Phi}} L_\Phi(g),$$

and these are by definition in the weighted L_2 space

$$L_{2,\sigma,1/\widehat{\Phi}}(D) := \left\{ u \ : \ \int_D \frac{|u(\omega)|^2}{\widehat{\Phi}(\omega)} d\sigma(\omega) < \infty \right\}.$$

5.2 Theory on the Torus using Fourier Series

5.3 Theory on Spheres using Expansions

5.4 Lower Bounds for Eigenvalues

(*SecLBE*) Here we proceed to prove lower bounds of the form (3.4.16, *GBound*) for the smallest eigenvalue of the matrix occurring in optimal recovery problems with Lagrange data. We had to postpone them until now, because they require transforms.

5.5 Generalizations of Results Using Transforms

6 Theory on Grids

Using Fourier transforms, we treat the case of gridded data $h\mathbb{Z}^d$ here.

6.1 Strang-Fix Theory

6.2 Application to Radial Basis Functions

6.3 Shift Invariant Spaces

7 Construction and Characterization of Positive Definite Functions

(*SecCCPD*) This section is intended to give the proofs of conditional positive definiteness of the classical radial basis functions. We include a toolbox of operators on radial functions that allow the construction of compactly supported positive definite functions.

7.1 General Construction Techniques

(*SecGCT*)

7.2 Construction of Radial Functions on $\mathbb{R}^d$

7.2.1 Gaussians

(*SecPDG*)

7.2.2 Nonexistence of CS Functions for All Dimensions

(*NECSAlld*)

7.3 Positive Definite Functions on Topological Groups

7.4 Positive Definite Zonal Functions on Spheres

8 Special Algorithms

(*SecSA*) This section contains some additional techniques that may be useful for the numerical solution of multivariate recovery problems.

8.1 Reduction of Enlarged System, Method 1

(*Red1*) Consider the enlarged system (1.7.3, *BDef*) and perform a partial Gaussian elimination algorithm on the matrix P with row interchanges. The result can be written in the form

$$LP\Pi = \begin{pmatrix} U \\ 0 \end{pmatrix}, \quad L = \begin{pmatrix} L_{11} & 0 \\ L_{21} & E \end{pmatrix}$$

with nonsingular lower triangular matrices L and L_{11} of size $M \times M$ and $q \times q$, respectively, with an $M \times M$ permutation matrix Π and a nonsingular $q \times q$ upper triangular matrix U, while L_{21} is some $M \times q$ matrix and E is the identity matrix. Now write α as a vector

(*split1*)

$$\alpha = L^T\beta = \begin{pmatrix} L_{11}^T & L_{21}^T \\ 0 & E \end{pmatrix} \begin{pmatrix} \beta^{(1)} \\ \beta^{(2)} \end{pmatrix} = \begin{pmatrix} L_{11}^T\beta^{(1)} + L_{21}^T\beta^{(2)} \\ \beta^{(2)} \end{pmatrix}, \tag{8.1.1}$$

again using a split of an M-vector into a q-vector followed by an $(M-q)$-vector. Ignoring the details of such obvious splits from now on, we evaluate

$$0 = \Pi^T P^T \alpha = \Pi^T P^T L^T \beta = U^T \beta^{(1)} + 0$$

and get $\beta^{(1)} = 0$. Now we split the system $LAL^T\beta + LP\gamma = Lf$ in the same way to get

$$\begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \begin{pmatrix} 0 \\ \beta^{(2)} \end{pmatrix} + \begin{pmatrix} U \\ 0 \end{pmatrix} \delta = \begin{pmatrix} g^{(1)} \\ g^{(2)} \end{pmatrix}$$

introducing the vector δ via $\gamma = \Pi\delta$. This decomposes into two systems

$$A_{22}\beta^{(2)} = g^{(2)}, \quad A_{12}\beta^{(2)} + U\delta = g^{(1)}$$

that can be solved for $\beta^{(2)}$ and δ, respectively. From these it is easy to calculate α and γ.

To see the positive definiteness of the matrix A_{22}, observe that

$$(\beta^{(2)})^T A_{22} \beta^{(2)} = \begin{pmatrix} 0 \\ \beta^{(2)} \end{pmatrix}^T LAL^T \begin{pmatrix} 0 \\ \beta^{(2)} \end{pmatrix}$$

holds for all $\beta^{(2)} \in I\!R^{M-q}$, and all α with (1.6.3, *CPDef*) have a unique split in the form (8.1.1, *split1*). Thus A_{22} defines a positive definite quadratic form on R^{M-q}, and it must be a positive definite matrix.

To calculate the numerical effort, we now explicitly write down the algorithm:

1. Perform q Gaussian transformations on rows of P with pivoting. This requires $\mathcal{O}(Mq^2)$ operations and generates the matrices Π, U, L_{11}, and L_{21}. The latter three can be stored over P, and Π requires an integer array of length q for keeping track of row permutations.

2. Calculate the submatrices A_{ik} of LAL^T by applying the Gaussian transformations stored in L to A row- and columnwise. These are q transformations of M-vectors each, and the overall effort will be $\mathcal{O}(M^2q)$. Note that this operation will cause fill-in, if the original matrix was sparse.

3. Calculate Lf and split it into $g^{(1)}$ and $g^{(2)}$. Using the special form of L again, this amounts to $\mathcal{O}(Mq)$ operations.

4. Solve the positive definite $(M-q)\times(M-q)$ system $A_{22}\beta^{(2)} = g^{(2)}$ by your favourite method. We shall comment on such problems for unconditionally positive definite functions later in section 2.3 (*CompEffort)*. Its computational complexity does not enter into the complexity of the transformation we consider here.

5. Now solve $A_{12}\beta^{(2)} + U\delta = g^{(1)}$ for δ. Using the upper triangular structure of U, the computational effort is $\mathcal{O}(Mq + q^2)$ for forming the system and solving it.

6. Backpermutation of elements of δ yields γ at $\mathcal{O}(q)$ cost.

7. Finally, α is an extension of $\beta^{(2)}$ by the q components of the vector $L_{21}^T\beta^{(2)}$, and these can be calculated by $\mathcal{O}(Mq^2)$ operations.

Since we started with a conditionally positive definite function of *positive* order m, the increase of Φ towards infinity leads to a matrix A that shows a more or less strong increase of elements when moving away from the main diagonal. After the reduction the resulting matrix behaves like one generated by a positive definite function (this is actually provable for the reduction method of the next section). Thus it usually shows off-diagonal decay, and numerical experiments indicate some improvement of the condition. Thus there is some hope that variations of these reduction methods can possibly be turned into efficient preconditioning techniques.

8.2 Reduction of Enlarged System, Method 2

(*Red2*) Again, we consider the enlarged system (1.7.3, *BDef*), but now we perform q Householder transformations on P^T with column pivoting. This

means a reordering of the points in $X = \{x_1, \ldots, x_M\}$ and transition to a new basis in $I\!P_m^d$. In linear algebra terms we end up with a decomposition

(*Dec2*)

$$U^{-1}QP^T\Pi = (E, S) \tag{8.2.1}$$

with a nonsingular upper triangular $q \times q$ matrix U, an orthogonal $q \times q$ matrix Q, an $M \times M$ permutation matrix Π and a plain $q \times (M-q)$ matrix S. Note that the Householder transformations first produce $QP^T\Pi = (U, *)$, but we left-multiply this with U^{-1} to get (8.2.1, *Dec2*).

Now we permute and split α by

(*split3*)

$$\alpha = \Pi\beta, \quad \beta = \begin{pmatrix} \beta^{(1)} \\ \beta^{(2)} \end{pmatrix} \tag{8.2.2}$$

into a q-vector followed by an $(M-q)$-vector. Then we evaluate

$$0 = U^{-1}QP^T\alpha = U^{-1}QP^T\Pi\Pi^{-1}\alpha = (E, S)\begin{pmatrix} \beta^{(1)} \\ \beta^{(2)} \end{pmatrix}$$

and get $\beta^{(1)} = -S\beta^{(2)}$. Now we split the system $\Pi^T A\Pi\beta + \Pi^T P\gamma = \Pi^T f$ to get

$$\begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}\begin{pmatrix} \beta^{(1)} \\ \beta^{(2)} \end{pmatrix} + \begin{pmatrix} E \\ S^T \end{pmatrix}\delta = \begin{pmatrix} g^{(1)} \\ g^{(2)} \end{pmatrix}$$

introducing the vector $\delta = U^TQ\gamma$. This decomposes into two systems

$$\begin{aligned} A_{11}\beta^{(1)} &+ A_{12}\beta^{(2)} &+ \delta &= g^{(1)} \\ A_{21}\beta^{(1)} &+ A_{22}\beta^{(2)} &+ S^T\delta &= g^{(2)} \end{aligned}$$

and we solve the first for δ by

(*DeltaSys*)

$$\delta = g^{(1)} - A_{11}\beta^{(1)} - A_{12}\beta^{(2)} = g^{(1)} + (A_{11}S - A_{12})\beta^{(2)}. \tag{8.2.3}$$

Putting this with $\beta^{(1)} = -S\beta^{(2)}$ into the second yields a symmetric $(M-q) \times (M-q)$ system

(*RedSys3*)

$$\left(A_{22} + S^TA_{11}S - S^TA_{12} - A_{21}S\right)\beta^{(2)} = g^{(2)} - S^Tg^{(1)} \tag{8.2.4}$$

that can be solved for $\beta^{(2)}$. To see the positive definiteness of the matrix (8.2.4, *RedSys3*), observe that

$$(\beta^{(2)})^T \left(A_{22} + S^T A_{11} S - S^T A_{12} - A_{21} S \right) \beta^{(2)} = \beta^T \Pi^T A \Pi \beta = \alpha^T A \alpha$$

holds for all $\beta^{(2)} \in I\!R^{M-q}$, and all α with (1.6.3, *CPDef*) have a unique split in the form (8.2.2, *split3*) with $\beta^{(1)} = -S\beta^{(2)}$. In 3.3.2 (*PhiNormalization)* we shall see that this matrix can be written in the form $A_{Y,\Psi}$ for a set Y of $M - q$ points and a function Ψ that is *unconditionally* positive definite on $I\!R^d \setminus (X \setminus Y)$..

Let us now explicitly write down the algorithm:

1. Perform q Householder transformations on P^T with pivoting by column permutation of P^T. This requires $\mathcal{O}(Mq^2)$ operations and generates the matrices Π, U, and US. The latter two can be stored over P^T, and Π requires an integer array of length M for keeping track of point permutations.

2. Solve for S by backward substitution, using U. This again requires $\mathcal{O}(Mq^2)$ operations and generates S, which can be stored over part of P.

3. Generate the submatrices A_{ik} of $\Pi^T A \Pi$ by applying the permutations defined by Π to A row- and columnwise. This requires $2M$ swaps of M-vectors, and the overall effort will be $\mathcal{O}(M^2)$. Note that this operation can be avoided by using indirect indexing in later steps, but be aware of the fact that indirect indexing spoils the positive effect of cache memory.

4. Permute the right-hand side of the system and split it into $g^{(1)}$ and $g^{(2)}$. This amounts to $\mathcal{O}(M)$ operations, but is unnecessary if indirect indexing is implemented.

5. The bulk of work in this reduction method lies in forming the positive definite matrix

$$A_{22} + S^T A_{11} S - S^T A_{12} - A_{21} S,$$

and it is of order $\mathcal{O}(M^2 q)$.

6. Now solve the positive definite $(M - q) \times (M - q)$ system (8.2.4, *RedSys3*) for $\beta^{(2)}$ by your favourite method. We considered such problems for unconditionally positive definite functions in section 2.3 (*CompEffort)*. Its computational complexity does not enter into the complexity of the transformation we describe here.

7. Now form $\beta^{(1)} = -S\beta^{(2)}$ with $\mathcal{O}(M^2 q)$ operations and

8. use (8.2.3, *DeltaSys*) to calculate δ with $\mathcal{O}(Mq^2)$ operations. The solution vector α just is a permuted version of β, but the calculation of γ requires solution of the system $U^T Q\gamma = \delta$ in two steps:

9. Calculate $Q\gamma$ from δ by backward substitution with $\mathcal{O}(q^2)$ operations, and

10. form $\gamma = Q^T(Q\gamma)$ by premultiplication of $Q\gamma$ with Q^T with $\mathcal{O}(q^3)$ operations. Since $M \geq q$ follows from $\mathbb{P}_m^d$-nongegeneracy of $X = \{x_1, \ldots, x_M\}$, this is at most an $\mathcal{O}(Mq^2)$ effort.

9 Computational Geometry Techniques

(*SecCGT*) This section contains algorithms from Computational Geometry that are useful for solving scattered data problems in the the large. The main topic will be the **k-nearest neighbor problem** and related **query problems**.

9.1 Voronoi Diagrams

(*SecVor*)

10 Appendix

10.1 Basis Functions

(*SecBF*) Here we try to give a complete list (up to this date) of the available conditionally positive definite functions with their transforms and their recursion formulas. Proofs are either in the main text or in section 10.4 (*SecSFT)* of the appendix.

10.2 MATLAB routines

Here we provide the MATLAB sources required to do the examples of section 2.5 (*SecExamples)*.

$\phi(r)$	Parameters	m
r^β	$\beta > 0,\ \beta \notin 2\mathbb{N}$	$m \geq \lceil \beta/2 \rceil$
$r^\beta \log r$	$\beta > 0,\ \beta \in 2\mathbb{N}$	$m > \beta/2$
$(r^2 + c^2)^{\beta/2}$	$\beta > 0,\ \beta \notin 2\mathbb{N}$	$m \geq \lceil \beta/2 \rceil$

Table 7: Conditionally Positive Definite Functions (*TCPDFct2*)

$\phi(r)$	Parameters	Smoothness	Dimension	Name/Reference
$e^{-\beta r^2}$	$\beta > 0$	$C^\infty(\mathbb{R}^d)$	$d < \infty$	Gaussian
$(r^2 + c^2)^{\beta/2}$	$\beta < 0$	$C^\infty(\mathbb{R}^d)$	$d < \infty$	inv. Multiquadric
$r^\nu K_\nu(r)$	$\nu > 0$	$C^{\lfloor \nu \rfloor}$	$d < \infty$	Sobolev spline
$(1-r)_+^2(2+r)$		C^0	$d \leq 3$	Wu [**?**](*wu:94-1*)
$(1-r)_+^4(1+4r)$		C^2	$d \leq 3$	Wendland [30](*wendland:95-1*)

Table 8: Unconditionally Positive Definite Functions (*TPDFct2*)

10.3 Hilbert Space Basics

(*SecHSB*) This is intended as a short tutorial on Hilbert spaces as required in Section 3 (*SecGT*). We only require fundamentals on linear spaces, bilinear forms, and norms. If a reader has problems with any of the stated facts below, it is time to go back to an introductory course on Calculus and Numerical Analysis.

Definition 10.3.1 *(DefPHS) A set $\mathcal{H}$ and a mapping $(\cdot,\cdot)_\mathcal{H} : \mathcal{H} \times \mathcal{H} \to \mathbb{R}$ form a* **pre-Hilbert space** *over $\mathbb{R}$, if the following holds:*

1. *$\mathcal{H}$ is a vector space over $\mathbb{R}$.*
2. *$(\cdot,\cdot)_\mathcal{H}$ is a symmetric positive definite bilinear form.*

A symmetric positive bilinear form as $(\cdot,\cdot)_\mathcal{H} : \mathcal{H} \times \mathcal{H} \to \mathbb{R}$ is often called an **inner product** on $\mathcal{H}$. Then

(*NormDef*)

$$\|x\|_\mathcal{H}^2 := (x,x)_\mathcal{H},\ x \in \mathcal{H} \tag{10.3.2}$$

defines a norm on $\mathcal{H}$, and we assume all readers to be familiar with this notion. Sometimes, we shall weaken the assumptions on $(\cdot,\cdot)_\mathcal{H}$ and only ask for symmetry and positive semidefiniteness. Even in this more general situation, we have the **Cauchy-Schwarz inequality**

$$|(u,v)_\mathcal{H}| \leq |u|_\mathcal{H} |v|_\mathcal{H}$$

for all $u, v \in \mathcal{H}$, where we use the notation $|x|^2_{\mathcal{H}} := (x,x)_{\mathcal{H}}$ to denote a **seminorm** instead of a norm as in (10.3.2, *NormDef*). To prove the Cauchy-Schwarz inequality as a warm-up, just consider the quadratic function

$$\varphi(t) := |u + tv|^2_{\mathcal{H}} = |u|^2_{\mathcal{H}} + 2t(u,v)_{\mathcal{H}} + t^2|v|^2_{\mathcal{H}}.$$

It must be nonnegative, and thus it has none or a double real zero. This property is satisfied for a general function $\varphi(t) = at^2 + 2bt + c$, iff $b^2 \leq ac$ holds. But this is the square of the Cauchy-Schwarz inequality.

For completeness, we recall some basics from normed linear spaces:

1. A sequence $\{u_n\}_{n\in\mathbb{N}} \subset \mathcal{N}$ of a normed linear space $\mathcal{N}$ with norm $\|\cdot\|_{\mathcal{N}}$ is a **zero sequence** in $\mathcal{N}$, if the sequence $\{\|u_n\|_{\mathcal{N}}\}_{n\in\mathbb{N}}$ converges to zero in $\mathbb{R}$.

2. A sequence $\{u_n\}_{n\in\mathbb{N}} \subset \mathcal{N}$ is a **convergent sequence** in $\mathcal{N}$ with limit u, if the sequence $\{u_n - u\}_n$ is a zero sequence.

3. A subspace $\mathcal{M}$ of $\mathcal{N}$ is a **closed subspace**, if for every convergent sequence $\{u_n\}_{n\in\mathbb{N}} \subset \mathcal{M} \subset \mathcal{N}$ with limit u one can conclude that the limit u also belongs to $\mathcal{M}$.

4. The normed linear space $\mathcal{N}$ is **complete** or a **Banach space**, if every sequence which is a Cauchy sequence in the norm $\|\cdot\|_{\mathcal{V}}$ is necessarily convergent in $\mathcal{V}$.

5. A complete normed linear space is closed, since each convergent sequence is a Cauchy sequence.

6. A subset $\mathcal{M}$ of a normed linear space $\mathcal{N}$ is **dense**, if each element of $\mathcal{N}$ can be written as a limit of a convergent sequence from $\mathcal{M}$.

Now we add some simple facts about pre-Hilbert spaces:

1. A mapping (or operator) $A : \mathcal{H} \to \mathcal{N}$ with values in a normed linear space $\mathcal{N}$ with norm $\|\cdot\|_{\mathcal{N}}$ is a **continuous mapping** or a **bounded mapping**, if there is a constant C such that

   $$\|Ax\|_{\mathcal{N}} \leq C\|x\|_{\mathcal{H}}$$

   holds for all $x \in \mathcal{H}$.

2. The mapping A then has an **operator norm**

$$\|A\|_{\mathcal{H},\mathcal{N}} := \sup_{x\in\mathcal{H}\setminus\{0\}} \frac{\|Ax\|_{\mathcal{N}}}{\|x\|_{\mathcal{H}}} \leq C$$

and the bound

$$\|Ax\|_{\mathcal{N}} \leq \|A\|_{\mathcal{H},\mathcal{N}}\|x\|_{\mathcal{H}}$$

is best possible.

3. Two subspaces $\mathcal{U},\ \mathcal{V}$ of a pre-Hilbert space are **orthogonal**, if all vectors $u \in \mathcal{U},\ v \in \mathcal{V}$ are orthogonal, i.e.: $(u, v)_{\mathcal{H}} = 0$.

Definition 10.3.3 *An element u^* of a subspace $\mathcal{M}$ of a normed linear space $\mathcal{N}$ is a* **best approximation** *to a given element $u \in \mathcal{N}$, if*

$$\|u - u^*\|_{\mathcal{N}} = \sup_{v\in\mathcal{M}} \|u - v\|_{\mathcal{N}} =: E_{\mathcal{M}}(u).$$

Theorem 10.3.4 *(BAT) An element u^* of a subspace $\mathcal{M}$ of a pre-Hilbert space $\mathcal{H}$ is a best approximation to a given element $u \in \mathcal{H}$, iff the variational identity*

(EqVar)

$$(u - u^*, v)_{\mathcal{H}} = 0 \text{ for all } v \in \mathcal{M} \tag{10.3.5}$$

holds. If it exists, the best approximation is unique. If $\mathcal{M}$ is finite-dimensional and spanned by linearly independent elements $u_1 \ldots, u_M$, then the coefficients α^* of the representation

$$u^* = \sum_{j=1}^{M} \alpha_j^* u_j$$

are solutions of the **normal equations**

$$\sum_{j=1}^{M} \alpha_j^* (u_j, u_k)_{\mathcal{H}} = (u, u_k)_{\mathcal{H}},\ 1 \leq k \leq M,$$

and the symmetric and positive definite matrix with entries $(u_j, u_k)_{\mathcal{H}}$ is called a **Gram matrix**.

Proof: Let u^* be a best approximation to u. Then consider an arbitrary $v \in \mathcal{M}$ and form the quadratic function

$$u_v(\alpha) := \|u - u^* + \alpha v\|_{\mathcal{H}}^2 = \|u - u^*\|_{\mathcal{H}}^2 + 2\alpha(u - u^*, v)_{\mathcal{H}} + \alpha^2\|v\|_{\mathcal{H}}^2$$

whose minimum must be attained at $\alpha = 0$. This implies $(u - u^*, v)_{\mathcal{H}} = 0$. Conversely, assume (10.3.5, *EqVar*) and write any other element $v \in \mathcal{M}$ as $v = u^* + 1 \cdot (v - u^*)$. Then (10.3.5, *EqVar*) implies that the quadratic function u_{u^*-v} is minimal at $\alpha = 0$, proving $u_{u^*-v}(1) = \|u - v\|_{\mathcal{H}} \geq u_{u^*-v}(0) = \|u - u^*\|_{\mathcal{H}}$. If u^* and u^{**} are two best approximations from $\mathcal{M}$ to u, then we can subtract the two variational identities $(u - u^*, v)_{\mathcal{H}} - (u - u^{**}, v)_{\mathcal{H}} = (u^{**} - u^*, v)_{\mathcal{H}} = 0$ for all $v \in \mathcal{M}$ and insert $v = u^{**} - u^*$ to get $u^{**} = u^*$. The third assertion is a specialization of (10.3.5, *EqVar*). □

Corollary 10.3.6 *The first statement of Theorem 10.3.4 (BAT) holds also in the case of a positive semidefinite bilinear form. The Gram matrix in the finite-dimensional case now is only positive semidefinite.* □

Corollary 10.3.7 *(BAC) Let $\lambda_1, \ldots, \lambda_M$ be linear functionals on a pre-Hilbert space $\mathcal{H}$ and let some $u \in \mathcal{H}$ be given. An element u^* of $\mathcal{H}$ solves the problem*

$$\|u^*\|_{\mathcal{H}} = \inf_{\substack{v \in \mathcal{H} \\ \lambda_j(v) = \lambda_j(u) \\ 1 \leq j \leq M}} \|v\|_{\mathcal{H}},$$

iff the variational identity

$$(u^*, v)_{\mathcal{H}} = 0 \text{ for all } v \in \mathcal{H} \text{ with } \lambda_j(v) = 0,\ 1 \leq j \leq M.$$

holds, or iff there are real numbers $\alpha_1, \ldots, \alpha_M$ such that

$$(u^*, v)_{\mathcal{H}} = \sum_{j=1}^{M} \alpha_j \lambda_j(v) \text{ for all } v \in \mathcal{H}.$$

Proof: Consider the subspace

$$\mathcal{M} = \{\, v \in \mathcal{H} \,:\, \lambda_j(v) = 0,\ 1 \leq j \leq M \}$$

and reformulate the problem by writing any $v \in \mathcal{H}$ with $\lambda_j(v) = \lambda_j(u),\ 1 \leq j \leq M$ as $v = u - w$ for $w \in \mathcal{M}$. Then we have a problem of best approximation to u from $\mathcal{M}$ and can simply use Theorem 10.3.4 (*BAT*) to prove the first assertion. We then have to prove that the first variational identity implies the second. But this follows from a standard linear algebra argument that we include for completeness as the next lemma. □

Lemma 10.3.8 *If $A : X \to Y$ and $B : X \to Z$ are linear maps between linear spaces, and if B vanishes on the kernel $\ker A$ of A, then B factorizes over $A(X)$, i.e.: there is a map $C : A(X) \to Z$ such that $B = C \circ A$. If Z is normed and if Y is finite-dimensional, then C is continuous.*

Proof: There is an isomorphism $D : A(X) \to X/\ker A$, and one can define $\widetilde{B} : A/\ker A \to Z$ by $\widetilde{B}(x + \ker A) := B(x)$ because $B(\ker A) = \{0\}$. Then $C := \widetilde{B} \circ D$ does the job, since

$$C(A(x)) = \widetilde{B}(D(A(x))) = \widetilde{B}(x + \ker A) = B(x)$$

for all $x \in X$. If Y is finite-dimensional, the isomorphic spaces $A(X) \subseteq Y$ and $X/\ker A$ must also be finite-dimensional. Since all linear mappings defined on finite-dimensional linear spaces with values in normed linear spaces are continuous, we are finished. □

So far, Theorem 10.3.4 *(BAT)* does not imply existence of best approximations from subspaces of infinite dimension. It just characterizes them. To get existence, we need that certain nice sequences actually have limits:

Definition 10.3.9 *(DefHS) A pre-Hilbert space $\mathcal{H}$ with inner product $(\cdot,\cdot)_{\mathcal{H}}$ is a* **Hilbert space** *over $\mathbb{R}$, if $\mathcal{H}$ is* **complete** *under the norm $\|\cdot\|_{\mathcal{H}}$, i.e.: as a normed linear space.*

We now prove the crucial **projection theorem** in Hilbert spaces:

Theorem 10.3.10 *(PTHS) If $\mathcal{H}$ is a Hilbert space with a closed subspace $\mathcal{M}$, then any element $u \in \mathcal{H}$ has a unique best approximation $u^*_{\mathcal{M}}$ from $\mathcal{M}$, and the elements $u^*_{\mathcal{M}}$ and $u - u^*_{\mathcal{M}}$ are orthogonal. The map $\Pi_{\mathcal{M}} : \mathcal{H} \to \mathcal{M}$ with $\Pi_{\mathcal{M}}(u) := u^*_{\mathcal{M}}$ is linear, has norm one if $\mathcal{M}$ is nonzero, and is a* **projector***, i.e.: $\Pi^2_{\mathcal{M}} = \Pi_{\mathcal{M}}$. If Id is the identity mapping, then $Id - \Pi_{\mathcal{M}}$ is another projector, mapping $\mathcal{H}$ onto the* **orthogonal complement**

$$\mathcal{M}^\perp := \{\, u \in \mathcal{H} \;:\; (u,v)_{\mathcal{H}} = 0 \text{ for all } v \in \mathcal{M} \,\}.$$

of $\mathcal{M}$. Finally, the decomposition

$$\mathcal{H} = \mathcal{M} + \mathcal{M}^\perp$$

is a direct and orthogonal sum of spaces.

Proof: The existence proof for approximations from finite-dimensional subspaces is a consequence of Theorem 10.3.4 *(BAT)*, and we postpone the

general case for a moment. The orthogonality statement follows in general from Theorem 10.3.4 (*BAT*), and it yields Pythagoras' theorem in the form

$$\|u\|_{\mathcal{H}}^2 = \|u - u_{\mathcal{M}}^*\|_{\mathcal{H}}^2 + \|u^*\|_{\mathcal{H}}^2.$$

This in turn proves that both projectors have a norm not exceeding one. It is easy to prove that $\alpha u_{\mathcal{M}}^* + \beta v_{\mathcal{M}}^*$ is a best approximation to $\alpha u + \beta v$ for all $\alpha, \beta \in I\!R$ and all $u, v \in \mathcal{H}$, using the variational identity in Theorem 10.3.4 (*BAT*). To prove linearity of the projectors, we use uniqueness of the best approximation, as follows from Theorem 10.3.4 (*BAT*). Finally, surjectivity of the projectors is easily proven from the best approximation property of their definition.

Thus we are left with the existence proof for the infinite-dimensional case. The nonnegative real number $E_{\mathcal{M}}(u)$ can be written as the limit of a decreasing sequence $\{\|u - v_n\|_{\mathcal{H}}\}_n$ for certain elements $v_n \in \mathcal{M}$, because it is defined as an infimum. Forming the subspaces

$$\mathcal{M}_n := \text{span}\,\{v_1, \ldots, v_n\} \subseteq \mathcal{M}$$

and unique best approximations w_n to u from $\mathcal{M}_n$, we get

$$E_{\mathcal{M}}(u) \leq \|u - w_n\|_{\mathcal{H}} \leq \|u - v_n\|_{\mathcal{H}},$$

such that the sequence $\{\|u - w_n\|_{\mathcal{H}}\}_n$ converges to $E_{\mathcal{M}}(u)$, too. We now fix indices $m \geq n$ and use that $(u - w_m, w_m - w_n)_{\mathcal{H}} = 0$ follows from the best approximation property of w_m. Then we have

$$\begin{aligned} \|u - w_n\|_{\mathcal{H}}^2 - \|u - w_m\|_{\mathcal{H}}^2 &= \|u - w_m + w_m - w_n\|_{\mathcal{H}}^2 - \|u - w_m\|_{\mathcal{H}}^2 \\ &= \|u - w_m\|_{\mathcal{H}}^2 + 2(u - w_m, w_m - w_n)_{\mathcal{H}} \\ &\quad + \|w_m - w_n\|_{\mathcal{H}}^2 - \|u - w_m\|_{\mathcal{H}}^2 \\ &= \|w_m - w_n\|_{\mathcal{H}}^2, \end{aligned}$$

and since the sequence $\{\|u - w_n\|_{\mathcal{H}}^2\}_n$ is convergent and thus a Cauchy sequence, we get that $\{w_n\}_n \subset \mathcal{M}$ is a Cauchy sequence in $\mathcal{M} \subseteq \mathcal{H}$. Now the completeness of $\mathcal{H}$ assures the existence of a limit $w^* \in \mathcal{H}$ of this sequence, and since $\mathcal{M}$ was ssumed to be closed, the element w^* must belong to $\mathcal{M}$. The above identity can be used to let m tend to infinity, and then we get

$$\|u - w_n\|_{\mathcal{H}}^2 - \|u - w^*\|_{\mathcal{H}}^2 = \|w^* - w_n\|_{\mathcal{H}}^2.$$

This proves

$$E_{\mathcal{M}}(u) \leq \|u - w^*\|_{\mathcal{H}} \leq \|u - w_n\|_{\mathcal{H}},$$

and since the right-hand side converges to $E_{\mathcal{M}}(u)$, the element w^* must be the best approximation to u. □

We now proceed towards the completion theorem for pre-Hilbert spaces.

Theorem 10.3.11 *(HSCT) Let $\mathcal{H}$ be a pre-Hilbert space with inner product $(\cdot,\cdot)_{\mathcal{H}}$. Then there is a Hilbert space $\mathcal{J}$ and an isometric embedding $J : \mathcal{H} \to \mathcal{J}$ such that the following is true:*

1. *$J(\mathcal{H})$ is dense in $\mathcal{J}$.*
2. *Any continuous mapping $A : \mathcal{H} \to \mathcal{N}$ with values in a Banach space $\mathcal{N}$ has a unique continuous extension $B : \mathcal{J} \to \mathcal{N}$ such that $B \circ J = A$.*

Proof: We first form the space of all Cauchy sequences in $\mathcal{H}$, which clearly is a linear space over $\mathbb{R}$. Two such sequences are called equivalent, if their difference is a sequence in $\mathcal{H}$ converging to zero. The space $\mathcal{J}$ now is defined as the space of equivalence classes of Cauchy sequences in $\mathcal{H}$ modulo zero sequences. These classes clearly form a vector space under the usual operations on sequences. If we use an overstrike to stand for "class of", we write an element of $\mathcal{J}$ as $\overline{\{u_n\}_n}$ for a Cauchy sequence $\{u_n\}_n \in \mathcal{H}$. Now it is time to define an inner product

$$(\overline{\{u_n\}_n}, \overline{\{v_n\}_n})_{\mathcal{J}} := \lim_{n\to\infty} (u_n, v_n)_{\mathcal{H}}$$

on $\mathcal{J}$ and the embedding J via the constant Cauchy sequences

$$Ju := \{u\}_n := \{u_n = u\}_n$$

for each $u \subset \mathcal{H}$. Then

$$(Ju, Jv)_{\mathcal{J}} = (u, v)_{\mathcal{H}}$$

makes sure that J is an isometry and injective. But we still have to show that the inner product on $\mathcal{J}$ is well-defined and positive definite. If $\{u_n\}_n$ and $\{v_n\}_n$ are Cauchy sequences in $\mathcal{H}$, then

$$|\|u_n\|_{\mathcal{H}} - \|u_m\|_{\mathcal{H}}| \le \|u_n - u_m\|_{\mathcal{H}}$$

implies that the sequences $\{\|u_n\|_{\mathcal{H}}\}_n$ and $\{\|v_n\|_{\mathcal{H}}\}_n$ are Cauchy sequences in $\mathbb{R}$, and thus convergent and bounded by constants C_u and C_v. But then

$$\begin{aligned}(u_n, v_n)_{\mathcal{H}} - (u_m, v_m)_{\mathcal{H}} &= (u_n, v_n)_{\mathcal{H}} - (u_n, v_m)_{\mathcal{H}} - (u_m, v_m)_{\mathcal{H}} + (u_n, v_m)_{\mathcal{H}} \\ &= (u_n, v_n - v_m)_{\mathcal{H}} - (u_m - u_n, v_m)_{\mathcal{H}} \\ &\le C_u \|v_n - v_m\|_{\mathcal{H}} + C_v \|u_m - u_n\|_{\mathcal{H}}\end{aligned}$$

proves that $\{(u_n, v_n)_{\mathcal{H}}\}_n$ is a Cauchy sequence in $\mathbb{R}$ and thus convergent. Two representatives of a class $\overline{\{u_n\}_n}$ differ just by a zero sequence that does not affect the inner product's value. The proof of definiteness again uses that zero sequences represent zero in $\mathcal{J}$. This finishes the proof of well-definedness of the new inner product.

Thus $\mathcal{J}$ is another pre-Hilbert space that contains an isometric image of $\mathcal{H}$, and we first want to prove that $J(\mathcal{H})$ is dense in $\mathcal{J}$. Let us take an element $\overline{\{u_n\}_n} \in \mathcal{J}$ and use the fact that for each $\epsilon > 0$ there is some $K(\epsilon)$ such that for all $n, m \geq K(\epsilon)$ we have

$$\|u_n - u_m\|_{\mathcal{H}} \leq \epsilon.$$

Now take $m \geq K(\epsilon)$ and the fixed Cauchy sequence $\overline{\{u_m\}_n} = J(u_m)$. Then

$$\|J(u_m) - \overline{\{u_n\}_n}\|_{\mathcal{J}} = \lim_{n\to\infty} \|u_m - u_n\|_{\mathcal{H}} \leq \epsilon$$

proves the density assertion.

We now proceed to prove completeness of $\mathcal{J}$. To do this we have to form a Cauchy sequence $\{\overline{\{u_n^{(m)}\}_n}\}_m$ of equivalence classes $\overline{\{u_n^{(m)}\}_n}$ of Cauchy sequences $\{u_n^{(m)}\}_n \subset \mathcal{H}$. For each $m \in I\!N$ we can use the density property of $\mathcal{H}$ in $\mathcal{J}$ to find an element $v_m \in \mathcal{H}$ such that

$$\|\overline{\{u_n^{(m)}\}_n} - J(v_m)\|_{\mathcal{J}} \leq 1/m.$$

Due to

$$\begin{aligned} \|v_n - v_m\|_{\mathcal{H}} &= \|J(v_n) - J(v_m)\|_{\mathcal{J}} \\ &\leq \|J(v_n) - \overline{\{u_n^{(n)}\}_n}\|_{\mathcal{J}} + \\ &\quad + \|\overline{\{u_n^{(n)}\}_n} - \overline{\{u_n^{(m)}\}_n}\|_{\mathcal{J}} + \|\overline{\{u_n^{(m)}\}_n} - J(v_m)\|_{\mathcal{J}} \\ &\to 0 \end{aligned}$$

for $n, m \to \infty$, the sequence $\{v_m\}_m$ is a Cauchy sequence in $\mathcal{H}$. We now form

$$\begin{aligned} \|\overline{\{u_n^{(k)}\}_n} - \{v_n\}_n\|_{\mathcal{J}} &\leq \|\overline{\{u_n^{(k)}\}_n} - J(v_k)\|_{\mathcal{J}} + \|J(v_k) - \{v_n\}_n\|_{\mathcal{J}} \\ &\leq 1/k + \lim_{n\to\infty} \|v_k - v_n\|_{\mathcal{H}} \\ &\to 0 \end{aligned}$$

for $k \to \infty$, proving convergence towards $\{v_n\}_n$.

Now let $A : \mathcal{H} \to \mathcal{N}$ be a linear continuous mapping with values in a complete normed linear space $\mathcal{N}$. If $\overline{\{u_n\}_n}$ is an element of $\mathcal{J}$, we define the extension B on $\overline{\{u_n\}_n}$ by

(*Bmapdef*)

$$B(\overline{\{u_n\}_n}) := \lim_{n\to\infty} A(u_n). \tag{10.3.12}$$

Since A is continuous, it is bounded and due to

$$\|A(u_m) - A(u_n)\|_{\mathcal{N}} \leq \|A\|\|u_m - u_n\|_{\mathcal{H}}$$

the sequence $\{Au_n\}_n$ is a Cauchy sequence in $\mathcal{N}$. But as $\mathcal{N}$ is a Banach space, the sequence is convergent and (10.3.12, *Bmapdef*) is well-defined. Clearly $B \circ J = A$ holds by definition. Any two such extensions must agree on the dense subspace $A(\mathcal{H})$ of $\mathcal{J}$, and since they are continuous, they must agree on all of $\mathcal{J}$. □We add a little application:

Lemma 10.3.13 *If $\mathcal{M}$ is a dense subspace of a Hilbert space $\mathcal{H}$, then the closure of $\mathcal{M}$ is isometrically isomorphic to $\mathcal{H}$.*

Proof: The closure of $\mathcal{M}$ can be identified with a closed subspace $\mathcal{N}$ of $\mathcal{H}$, and we assert that $\mathcal{N} = \mathcal{H}$. To this end, decompose $\mathcal{H}$ into $\mathcal{H} = \mathcal{N} + \mathcal{N}^\perp$ and take an element u from $\mathcal{N}^\perp$. It must be orthogonal to all elements from $\mathcal{M}$, and by continuity of the functional $v \mapsto (u, v)_{\mathcal{H}}$ it must be orthogonal to all of $\mathcal{H}$. Thus it must be zero. □

We further need the **Riesz representation theorem** for continuous linear functionals:

Theorem 10.3.14 *(RieszT) Any continuous linear real-valued functional λ on a Hilbert space $\mathcal{H}$ can be written as*

(RieszRep)

$$\lambda = (\cdot, g_\lambda)_{\mathcal{H}} \tag{10.3.15}$$

with a unique element $g_\lambda \in \mathcal{H}$. The map $\lambda \mapsto g_\lambda$ is an isometric isomorphism between the **dual Hilbert space** *$\mathcal{H}^*$ of $\mathcal{H}$, consisting of all continuous linear real-valued functionals on $\mathcal{H}$, and $\mathcal{H}$ itself.*

Proof: If $\lambda = 0$, then $g_\lambda = 0$ does the job and is unique. If $\lambda \neq 0$, the kernel $\mathcal{L}$ of λ is not the full space $\mathcal{H}$. It is, however, a closed linear subspace, and thus there is some element $h_\lambda \in \mathcal{L}^\perp$ with $\|h_\lambda\|_{\mathcal{H}} = 1$. Now for each $u \in \mathcal{H}$ the element $\lambda(u)h_\lambda - \lambda(h_\lambda)u$ must necessarily be in $\mathcal{L}$ and thus orthogonal to h_λ. This means

$$\begin{aligned} 0 &= (h_\lambda, \lambda(u)h_\lambda - \lambda(h_\lambda)u)_{\mathcal{H}}, \\ \lambda(u)(h_\lambda, h_\lambda)_{\mathcal{H}} &= \lambda(h_\lambda)(u, h_\lambda)_{\mathcal{H}}, \\ \lambda(u) &= (u, \lambda(h_\lambda)h_\lambda)_{\mathcal{H}}. \end{aligned}$$

The norm of λ is bounded by

$$\begin{aligned} \|\lambda\|_{\mathcal{H}^*} &:= \sup_{u \in \mathcal{H}\setminus\{0\}} \frac{|\lambda(u)|}{\|u\|_{\mathcal{H}}} \\ &\leq |\lambda(h_\lambda)| \end{aligned}$$

due to Cauchy-Schwarz, but using $u = h_\lambda$ in the definition of the norm yields equality. Since we set $g_\lambda := \lambda(h_\lambda)h_\lambda$, we get $\|\lambda\|_{\mathcal{H}^*} = \|g_\lambda\|_{\mathcal{H}}$. Uniqueness of g_λ satisfying (10.3.15, *RieszRep*) is easy to prove, and equally easy is the proof of injectivity and surjectivity of the map $\lambda \mapsto g_\lambda$. □

10.4 Special Functions and Transforms

(*SecSFT*) This is intended as a reference and tutorial for classical formulas involving special functions (e.g.: Bessel functions) and their transforms.

The **Gamma function** is defined by

(*GammaDef*)

$$\Gamma(z) = \int_0^\infty t^{z-1}e^{-t}dt. \tag{10.4.1}$$

The surface area of the $d-1$-dimensional sphere in $\mathbb{R}^d$ for $d > 1$ is

(*VolS*)

$$\mathrm{vol}\,(S^{d-1}) = 2\pi^{(d-1)/2}/\Gamma((d-1)/2). \tag{10.4.2}$$

The Bessel function K_ν of third kind (alias Mcdonald function) is

(*KnuDef*)

$$K_\nu(z) = \frac{\pi^{1/2}(z/2)^\nu}{\Gamma(\nu+1/2)}\int_1^\infty e^{-zt}(t^2-1)^{\nu-1/2}dt \tag{10.4.3}$$

for $|\arg z| < \pi/2$ and $\Re\nu > -1/2$, and its asymptotics near zero is

(*KnuAsyZero*)

$$K_\nu(z) = \frac{(z/2)^{-\nu}}{\Gamma(\nu)} + o(1), \tag{10.4.4}$$

while it behaves like

(*KnuAsyInf*)

$$K_\nu(z) = \frac{\sqrt{\pi}}{\sqrt{2z}}e^{-z}(1+\mathcal{O}(z^{-1})), \tag{10.4.5}$$

near infinity.

10.5 Necessary Results from Real Analysis

Here we collect some of the basic material on Lebesgue integration, Sobolev spaces, distributions, pseudodifferential operators, and partial differential equations.

10.5.1 Lebesgue Integration

(*SecLI*)

10.5.2 L_2 spaces

Lemma 10.5.1 *(LemContShift) The shift operator $S_z : f(\cdot) \mapsto f(\cdot - z)$ is a continuous function of z near zero in the following sense: for each given $u \in L_2(\mathbb{R}^d)$ and each given $\epsilon > 0$ there is some $\delta > 0$ such that*

$$\|S_z(u) - u\|_{L_2(\mathbb{R}^d)} \leq \epsilon$$

for all $\|z\|_2 \leq \delta$.

Proof: to be supplied later....

We now want to prove that the space $\mathcal{S}$ of tempered test functions is dense in $L_2(\mathbb{R}^d)$. For this, we have to study functions like (4.1.5, *deltaschar*) in some more detail. They are in $\mathcal{S}$ for all positive values of ϵ, and Lemma 4.1.6 (*LemRepro)* tells us that the operation

$$f \mapsto M_\epsilon(f) := \int_{\mathbb{R}^d} f(y)\varphi(\epsilon, \cdot - y)dy$$

maps each continuous L_1 function f to a "mollified" function $M_\epsilon(f)$ such that

$$\lim_{\epsilon \to 0} M_\epsilon(f)(x) = f(x)$$

uniformly on compact subsets of $\mathbb{R}^d$.

It is common to replace the Gaussian in (4.1.8, *deltarep*) by an infinitely differentiable function with compact support, e.g.

(*Friedmoll*)

$$\varphi(\epsilon, x) = \left\{ \begin{array}{ll} c(\epsilon)\exp(-1/(\epsilon^2 - \|x\|_2^2)) & \|x\|_2 < \epsilon \\ 0 & \|x\|_2 \geq \epsilon \end{array} \right\} \tag{10.5.2}$$

where the constant $c(\epsilon)$ is such that

$$\int_{\mathbb{R}^d} \varphi(\epsilon, x)dx = 1$$

holds for all $\epsilon > 0$. This **Friedrich's mollifier** can also be used in the definition of M_ϵ. It has the advantage that Lemma 4.1.6 (*LemRepro)* holds for more general functions, i.e.: for functions which are in L_1 only locally.

We now want to study the action of M_ϵ on L_2 functions. Let $u \in L_2(\mathbb{R}^d)$ be given, and apply the Cauchy-Schwarz inequality to

$$M_\epsilon(f)(x) = \int_{\mathbb{R}^d} (f(y)\sqrt{\varphi(\epsilon, x-y)})\sqrt{\varphi(\epsilon, x-y)}dy$$

to get

$$\begin{array}{rcl} |M_\epsilon(f)(x)|^2 & \leq & \int_{\mathbb{R}^d} |f(y)|^2 \varphi(\epsilon, x-y)dy \int_{\mathbb{R}^d} \varphi(\epsilon, x-y)dy \\ & = & \int_{\mathbb{R}^d} |f(y)|^2 \varphi(\epsilon, x-y)dy \end{array}$$

and

$$\int_{\mathbb{R}^d} |M_\epsilon(f)(x)|^2 dx \leq \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} |f(y)|^2 \varphi(\epsilon, z)dydz = \int_{\mathbb{R}^d} |f(y)|^2 dy$$

such that M_ϵ has norm less than or equal to one in the L_2 norm. It is even more simple to prove the identity

$$(f, M_\epsilon g)_{L_2(R^d)} = (M_\epsilon f, g)_{L_2(R^d)}$$

for all $f, g \in L_2(\mathbb{R}^d)$ by looking at the integrals. Here, we used the Fubini theorem on $\mathbb{R}^d$ which requires some care, but there are no problems because everything can either be done with a Friedrich's mollifier, or be done on sufficiently large compact sets first, letting the sets tend to $\mathbb{R}^d$ later.

We now use a Friedrich's mollifier to study the L_2 error of the mollification. This is very similar to the arguments we already know. The error is representable pointwise as

$$f(x) - M_\epsilon(f)(x) = \int_{\mathbb{R}^d} (f(x) - f(y))\varphi(\epsilon, x-y)dy$$

and we can use the Cauchy-Schwarz inequality to get

$$|f(x) - M_\epsilon(f)(x)|^2 \leq \int_{\|x-y\|_2 < \epsilon} |f(x) - f(y)|^2 \varphi(\epsilon, x-y)dy.$$

This can be integrated to get

$$\int_{\mathbb{R}^d} |f(x) - M_\epsilon(f)(x)|^2 dx \leq \int_{\|z\|_2 < \epsilon} \varphi(\epsilon, z) \int_{\mathbb{R}^d} |f(y+z) - f(y)|^2 dydz,$$

and we use the continuity of the shift operator as proven in Lemma 10.5.1 (*LemContShift)* to make this as small as we want by picking a suitably small ϵ. This shows

$$\lim_{\epsilon \to 0} \|f - M_\epsilon(f)\|_{L_2(\mathbb{R}^d)} = 0$$

and proves

Lemma 10.5.3 *(FTD) The space $\mathcal{S}$ of test functions is dense in $L_2(\mathbb{R}^d)$.*
□

10.5.3 Sobolev Spaces

(*SecSob*) This section contains definitions of Sobolev spaces and proves Sobolev's embedding theorems.

References

[1] J.H. Ahlberg, E.N. Nilson, and J.L. Walsh. *The theory of splines and their applications*, volume 38 of *Mathemtics in science and engineering*. Academic Press, 1986.

[2] M. Atteia. *Hilbertian Kernels and Spline Functions*. North-Holland, Amsterdam, 1992. Studies in Computational Mathematics (4).

[3] R.K. Beatson and G.N. Newsam. Fast evaluation of radial basis functions: I. Advances in the theory and applications of radial basis functions. *Comput. Math. Appl.*, 24(12):7–19, 1992.

[4] B. Bojanov and H. Woźniakowski, editors. *Optimal Recovery*. Nova Science Publishers, New York, 1992.

[5] C. de Boor, R.A. DeVore, and A. Ron. The structure of finitely generated shift–invariant spaces in $L_2(I\!R^d)$. *J. Funct. Anal.*, pages 37–78, 1994.

[6] J. Duchon. Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *Rev. Française Automat. Informat. Rech. Opér. Anal. Numer.*, 10:5–12, 1976.

[7] J. Duchon. Sur l'erreur d'interpolation des fonctions de plusieurs variables pas les D^m–splines. *Rev. Française Automat. Informat. Rech. Opér. Anal. Numer.*, 12:325–334, 1978.

[8] J. Duchon. Splines minimizing rotation–invariate semi–norms in sobolev spaces. In W. Schempp and K. Zeller, editors, *Constructive Theory of Functions of Several Variables*, pages 85–100. Springer, Berlin–Heidelberg, 1979.

[9] M.S. Floater and A. Iske. Multistep scattered data interpolation using compactly supported radial basis functions. To appear in *Journal of Computational and Applied Mathematics*, 1995.

[10] M.S. Floater and A. Iske. Scattered data analysis and industrial applications. In F. Fontanella, K. Jetter, and P.-J. Laurent, editors, *Advanced*

Topics in Multivariate Approximation, pages 1–10. World Scientific Publishing, 1996.

[11] R. Franke. Scattered data interpolation. Test of some methods. *Mathematics of Computation*, pages 181–200, 1982.

[12] M. Golomb and H.F. Weinberger. Optimal approximation and error bounds. In R.E. Langer, editor, *On Numerical Approximation*, pages 117–190. The University of Wisconsin Press, Madison, 1959.

[13] R.L. Hardy. Multiquadric equations of topography and other irregular surfaces. *J. Geophys. Res.*, 76:1905–1915, 1971.

[14] W.R. Madych and S.A. Nelson. Multivariate interpolation: a variational theory. Manuscript, 1983.

[15] W.R. Madych and S.A. Nelson. Multivariate interpolation and conditionally positive definite functions. *Approximation Theory and its Applications*, 4:77–89, 1988.

[16] W.R. Madych and S.A. Nelson. Error bounds for multiquadric interpolation. In C.K. Chui, L.L. Schumaker, and J.D. Ward, editors, *Approximation Theory VI*, pages 413–416. Academic Press, Boston, 1989.

[17] W.R. Madych and S.A. Nelson. Multivariate interpolation and conditionally positive definite functions II. *Mathematics of Computation*, 54:211–230, 1990.

[18] W.R. Madych and S.A. Nelson. Bounds on multivariate polynomials and exponential error estimates for multiquadric interpolation. *Journal of Approximation Theory*, 70:94–114, 1992.

[19] C.A. Micchelli. Interpolation of scattered data: distance matrices and conditionally positive definite functions. *Constructive Approximation*, 2:11–22, 1986.

[20] C.A. Micchelli and T.J. Rivlin. A survey of optimal recovery. In C.A. Micchelli and T.J. Rivlin, editors, *Optimal Estimation in Approximation Theory*, pages 1–54. Plenum Press, 1977.

[21] C.A. Micchelli and T.J. Rivlin. Optimal recovery of best approximations. *Resultate der Mathematik*, 3:25–32, 1978.

[22] C.A. Micchelli and T.J. Rivlin. Lectures on optimal recovery. In P.R. Turner, editor, *Numerical Analysis, Lancaster 1984*, pages 12–93. Springer Verlag, Lecture Notes in Mathematics 1129, 1984.

[23] C.A. Micchelli, T.J. Rivlin, and S. Winograd. Optimal recovery of smooth functions approximations. *Numerische Mathematik*, 260:191–200, 1976.

[24] I.P. Natanson. *Konstruktive Funktionentheorie*. Akademie–Verlag, Berlin, 1955.

[25] M.J.D. Powell. The theory of radial basis function approximation in 1990. In W.A. Light, editor, *Advances in Numerical Analysis II: Wavelets, Subdivision Algorithms, and Radial Basis Functions*, pages 105–210. Oxford Univ. Press, Oxford, 1992.

[26] M.J.D. Powell. Univariate multiquadric approximation: reproduction of linear polynomials. Technical Report DAMTP 1990/NA2, Department of Applied Mathematics and Theoretical Physics, University of Cambridge, 1992.

[27] M.J.D. Powell. Truncated Laurent expansions for the fast evaluation of thin plate splines. *Numer. Algorithms*, 5(1–4):99–120, 1993. Algorithms for Approximation, III (Oxford, 1992).

[28] R. Schaback. Error estimates and condition numbers for radial basis function interpolation. *Advances in Computational Mathematics*, 3:251–264, 1995.

[29] J. Stewart. Positive definite functions and generalizations, an historical survey. *Rocky Mountain J. Math.*, 6:409–434, 1976.

[30] H. Wendland. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in Computational Mathematics*, 4:389–396, 1995.

These were the references from the usual radial basis function database kept in Göttingen. The following are preliminary, and will be updated later.

References

[1] (*BraessBuch*) D. Braess, Nonlinear Approximation Theory, Springer 1986

[2] (*cheney:??-1*) E.W. Cheney, book on approximation theory

[3] (*meinardus:??-1*) G. Meinardus, book on approximation theory

[4] (*Meschkowski*) Meschkowski, Hilberträume mit Kernfunktion....

[5] (*natanson:55-1*) Natanson, Konstruktive Funktionentheorie,

[6] (*PowellEffTPSSystem*) Powell, efficiency of tps system enhanced

[7] (*NRSW*) Narcowich, Schaback, Ward, new paper multilevel

[8] (*RSTranslInv*) R.S., Translation Invariant...

[9] (*RSImpEB*) R.S., Improved Error Bounds

Index

This is a preliminary version: no subitems, just some keywords as provided by the `\kw` macro.

k-nearest neighbor problem 110
Banach space 112
Beppo-Levi space 45
Cauchy-Schwarz inequality 111
Euclidean norm 4
Fourier Inversion Theorem 95
Fourier transform 93
Friedrich's mollifier 121
Gamma function 120
Gaussian 10
Gaussian 93
Gram matrix 113
Hilbert space 115
Lipschitz continuity 75
Parseval's equation 103
Riesz representation theorem 119
Sobolev inequality 45
Sobolev space 46
Stirling's inequality 82
approximation 4
barycentric coordinates 80
best approximation 113
bounded mapping 112
closed subspace 112
complete 112
complete 115
condition number 66
conditionally positive definite function 11
conditionally positive definite function 42
continuous mapping 112
convergent sequence 112
dense 112
dual Hilbert space 119
fill distance 15
finite elements 7
finite interpolation property 39
finite separation property 39
generalized optimal power function 37
generalized optimal recovery problem 27
generalized transform 104
harmonic functions 46
inner product 111
interior cone condition 86
interpolation 4
inverse multiquadric 10
locations 6
multiquadric 10
native space 51
native space 58
nondegenerate 6
normal equations 113
normalization 58
operator norm 113
optimal recovery problem 26
orthogonal complement 115
orthogonal 113
polyharmonic functions 46
positive definite function 39
positive definite function 9
power function 62
pre-Hilbert space 111
projection theorem 115
projector 115
query problems 110
radial basis functions 7
recovery problem 4
rotational invariance 7
seminorm 112

separation distance 15
sphere 4
tempered test functions 93
test functions 102
thin-plate splines 46
thin-plate spline 10
transform mapping 102
translation invariance 7
unit ball 4
zero sequence 112
zonal 49
zonal 7
almost everywhere 102