Georg-August-Universität Göttingen

Convergence of Unsymmetric Kernel–Based Meshless Collocation Methods

Robert Schaback

Preprint Nr. 2005-07

Preprint-Serie des
Instituts für Numerische und Angewandte Mathematik
Lotzestr. 16-18
D - 37083 Göttingen

CONVERGENCE OF UNSYMMETRIC KERNEL–BASED MESHLESS COLLOCATION METHODS

ROBERT SCHABACK*

Abstract. This paper rigorously proves convergence of variations of the unsymmetric kernel-based collocation method introduced by E. Kansa in 1986. Since then, this method has been very successfully used in many applications, though it may theoretically fail in special situations, and though it had no error bound or convergence proof up to now. Thus it is necessary to add assumptions or to make modifications. Our modifications will prevent numerical failure and allow a rigorous mathematical analysis proving error bounds and convergence rates. These rates improve with the smoothness of the solution, the domain, and the kernel providing the trial spaces. More precisely, they are rates of approximation to the residuals by nonstationary meshless-kernel-based trial spaces. In particular, the rates are independent of the type of differential operator. The algorithms are applicable to large classes of linear problems in strong form, provided that there is a smooth solution and the test and trial discretizations are chosen with some care. Our theory does not require assumptions like ellipticity, and it can even handle ill-posed problems. Some numerical examples are added for illustration.

Key words. Kansa method, error bounds, stability, ill-posed problems, greedy adaptive solver

AMS subject classifications. 65N12, 65N35, 65N22, 65F22

1. Introduction. The final goal of this paper is to prove error bounds and convergence of certain numerical techniques that approximately solve a PDE problem via an unsymmetric or even non-square system of linear collocation equations. The most popular method of this kind was first proposed by E. Kansa [11] in 1986, and there are many follow-up papers in engineering journals (see e.g. [5] for a selection) that can easily be retrieved via the Internet. So far, the method is quite successful in applications with smooth solutions, but it can fail [10] in specially constructed situations. Consequently, it has neither error bounds nor convergence proofs for its original form, and a rigorous mathematical analysis will either require some additional assumptions or make changes to the method itself. We shall do both, but we shall stay general enough not to spoil the applicability to elliptic, parabolic, and hyperbolic problems. Therefore we need a somewhat nonstandard framework which we sketch here first, to make sure that the reader does not get lost in the technical details we have to provide later.

Consider a linear operator equation

$$L(u) = f,\ L\ :\ U \to F \tag{1.1}$$

between normed linear spaces U and F which is to be solved for any given $f \in F$. The map L takes a solution $u \in U$ to its *data* $L(u)$ in F. Thus F will usually be a Cartesian product of *trace* spaces of functions prescribed on the domain or on parts of the boundary. We shall consider a large class of unsymmetric discretization methods to solve such equations approximately, and we need five essential ingredients.

The first ingredient requires the problem to be *continuously dependent on the data*. In quantitative form, this means that the inverse L^{-1} is a bounded linear map from F to U. In particular, we assume an inequality of the form

$$\|u\|_U \le C_a \|L(u)\|_F \text{ for all } u \in U \tag{1.2}$$

*Institut für Numerische und Angewandte Mathematik, Lotzestraße 16–18, D–37083 Göttingen, Germany, http://www.num.math.uni-goettingen.de/schaback

with a positive constant C_a describing the stability of the problem. In practical applications this will imply that the *problem spaces* $U$and F have to be chosen with some care. In particular, U and F must often be chosen on a theoretical basis, e.g. as quite large spaces in which certain general existence results hold and which carry only rather weak norms. Usually, U will be a Sobolev space $W_2^\mu(\Omega)$ while F is a Cartesian product of Sobolev spaces that provide the right-hand sides for the differential equation and the boundary data via trace spaces. Continuous dependence serves here as a replacement for more specific analytic assumptions like coercivity of a bilinear form. However, in §9 we shall abandon the assumption of continuous dependence to be able to treat a certain class of ill-posed problems.

The second ingredient is some *additional regularity*. The actual solution u of a specific problem will often have more regularity than needed for the spaces U and F defining continuous dependence, and therefore we shall focus on a subspace $U_R \subseteq W_2^m(\Omega) \subset W_2^\mu(\Omega) =: U$ of U which we call the *regularity* subspace. The additional regularity of order $m - \mu > 0$ will be the driving force behind convergence rates, as we shall prove later.

The third ingredient is a scale of finite-dimensional *trial* subspaces U_r of U for a *trial discretization parameter* $r > 0$ which uses the additional regularity to provide a *convergent scheme for data approximation*. This is formalized by not necessarily linear maps $I_r : U_R \to U_r$ with error bounds

$$\|L(u - I_r(u))\|_F \le \epsilon_r(u) \text{ for all } u \in U_R. \tag{1.3}$$

It will be this approximation property that yields our convergence rates while feeding on the additional regularity. Note that we do not use a single discretization parameter like the usual h here, because we need two different scaling parameters r and s for trial and test discretization. Note further that we do not approximate the solution, but rather the data.

The fourth ingredient is a scale of *stable test discretizations* F_s of the data space F with respect to the scale of trial spaces U_r. This is formalized by a *test discretization parameter* $s > 0$ and linear maps $\Pi_s : F \to F_s$ into a scale of finite-dimensional spaces F_s such that the inequalities

$$\begin{aligned} \|L(u_r)\|_F &\le C(r,s)\|\Pi_s L(u_r)\|_{F_s} \text{ for all } u_r \in U_r \\ c(s)\|\Pi_s L(u)\|_{F_s} &\le \|L(u)\|_F \text{ for all } u \in U \end{aligned} \tag{1.4}$$

hold. We call a specific choice of trial and test discretization schemes *uniformly stable*, if both constants can be chosen independent of r and s. When restricted to the finite-dimensional test spaces U_r, the inequalities express equivalence of discrete and non-discrete norms. The second inequality will be easily satisfied by discretization, but the first one will be hard, because it bounds a non-discrete norm by a discrete norm, and this can work only for finite-dimensional spaces. It also implies uniqueness of solutions of the discretized finite systems $\Pi_s L(u_r) = \Pi_s L(u)$, which is a serious problem.

The final ingredient is the class of *numerical methods*. We do not specify details in this overview, but we can assume that we can always find elements $u^*_{r,s} \in U_r$ with

$$\|\Pi_s L(u - u^*_{r,s})\|_{F_s} \le \delta_{r,s}, \tag{1.5}$$

while these need not be unique. In fact, the approximation $I_r(u)$ is a solution, if we have

$$c(s)\epsilon_r(u) \le \delta_{r,s}. \tag{1.6}$$

Note that we do not attempt to solve the system $\Pi_s L(u_r) = \Pi_s L(u)$ exactly, because it will be overdetermined and unsolvable in general. However, under the assumption (1.6) we know that the relaxed problem (1.5) is solvable. In §8 we show how to solve such problems, and examples will follow in §10. Now we can formulate the core result of this paper:

THEOREM 1.1. *If the analytic problem is solvable by $u \in U_R$ and if we solve* (1.5) *by some $u^*_{r,s} \in U_r$, then there is an error bound*

$$\|u - u^*_{r,s}\|_U \leq C_a \left(\epsilon_r(u) \left(1 + \frac{C(r,s)}{c(s)} \right) + c(s)\delta_{r,s} \right).$$

If the discretization is uniformly stable, then there is a choice of $\delta_{r,s}$ via (1.6) *such that the above error behaves asymptotically like the trial approximation error $\epsilon_r(u)$.*

Proof. The assertion follows from a simple chain of inequalities:

$$\begin{aligned}
\tfrac{1}{C_a}\|u - u^*_{r,s}\|_U &\leq \|L(u - u^*_{r,s})\|_F \\
&\leq \|L(u - I_r(u))\|_F + \|L(I_r(u) - u^*_{r,s})\|_F \\
&\leq \epsilon_r(u) + c(s)\|\Pi_s L(I_r(u) - u^*_{r,s})\|_{F_s} \\
&\leq \epsilon_r(u) + c(s)\|\Pi_s L(u - u^*_{r,s})\|_{F_s} \\
&\quad + c(s)\|\Pi_s L(I_r(u) - u)\|_{F_s} \\
&\leq \epsilon_r(u) + c(s)\delta_{r,s} + \tfrac{C(r,s)}{c(s)}\|L(I_r(u) - u)\|_F \\
&\leq \epsilon_r(u)\left(1 + \tfrac{C(r,s)}{c(s)}\right) + c(s)\delta_{r,s}
\end{aligned}$$

□

But now we shall have to show how this abstract machinery can be set to work. We shall finally derive specific convergence rates for problems with continuous dependence in Sobolev spaces, including the Poisson problem with Dirichlet data as an illustration.

The following sections will treat the above ingredients one by one, and then we shall patch the results together. Our key tools will be *nonstationary meshless kernel-based trial spaces* which allow approximation schemes with high-order convergence rates while maintaining stability if paired with sufficiently rich test discretizations. It turns out that the use of smooth kernels makes the final convergence order dependent only on the regularity of the solution and the problem. The numerical methods for solving (1.5) will consist of certain variations of the original unsymmetric collocation method, and we already have solvability via (1.6). This saves us from the degeneration problems of the standard unsymmetric collocation technique [10]. A simple numerical example will finally illustrate these numerical techniques.

2. Well–Posed Problems and Regularity. For example, consider a standard Poisson boundary value problem

$$\begin{aligned}
-\Delta u &= f_\Omega \quad \text{in } \Omega \\
u &= f_D \quad \text{on } \partial\Omega
\end{aligned} \tag{2.1}$$

on a bounded domain $\Omega \subset \mathbb{R}^d$ with Dirichlet data f_D on the piecewise smooth boundary $\partial\Omega$. In such problems, we consider the equations as being given in strong form, i.e. we assume the solution u to be regular enough to pose the equations pointwise as

$$\begin{aligned}
(-\Delta u)(x) &= (\delta_x \circ (-\Delta))(u) = f_\Omega(x) \quad \text{for all } x \in \Omega \\
u(x) &= (\delta_x \circ Id)(u) = f_D(x) \quad \text{for all } x \in \partial\Omega.
\end{aligned} \tag{2.2}$$

But we allow much more general linear equations and boundary value operators. Formally, we follow the notation of [3] and others to combine differential and boundary operators into just one equation and write the latter as

$$L(u) \quad = \quad f \in F \tag{2.3}$$

where u is a function from some normed space U of functions. The mapping $L : U \to F$ maps solutions $u \in U$ to their data $L(u) \in F$, and the given problem consists in the inversion of L.

However, we want to solve the equation $L(u) = f$ in *strong* form, i.e. we interpret both sides as vectors of functions and want them to coincide pointwise. Note that this is still general enough to cover wide classes of ODE or PDE problems, but it requires regularity assumptions on the solution u which are stronger than the minimal smoothness properties guaranteed for solutions of weakly formulated problems. However, it is well-known that problems with regular data and well-shaped domains will have regular solutions.

When aiming at methods with error bounds and convergence, we have to take a closer look at the given analytic problem (2.3). In particular, we shall assume that the problem (2.3) is well–posed in the sense that the solution u depends continuously on the data f of the right–hand side of (2.3). But we have to make this more precise. This can be done in various ways, e.g. by *total sets* of data functionals, but this is not quantitative. For later use we impose a norm $\|.\|_F$ on $F := L(U)$ in a suitable way and assume (1.2) with a positive "analytic" constant C_a which describes the norm of the linear map L^{-1} that takes the data $f \in F$ and maps them back to the solution u in the function space U. Clearly, for such a–priori inequalities we must be careful with the choice of norms, because they depend on regularity theory, and they always imply that the homogeneous equation has only the trivial solution. The numerical methods following below will work on the discretized versions F_s, and thus the proper choice of F will also have practical consequences.

So far we have not mentioned any specific numerical algorithm. But if any numerical method has produced an approximate solution $\tilde{u} \in U$ to the problem (2.3), one can calculate the data $\tilde{f} = L(\tilde{u}) \in F$ and the norm $\|\tilde{f} - f\|_F$ to get the *a-posteriori* error bound

$$\|u - \tilde{u}\|_U \leq C_a \cdot \|L(u - \tilde{u})\|_F = C_a \cdot \|f - \tilde{f}\|_F \tag{2.4}$$

for free, since the *residuals* $L(u - \tilde{u}) = f - \tilde{f}$ are explicitly known. It means that errors in the solution are bounded by the norm of the residuals, multiplied with the analytic constant. Then the following simple observation is very useful for assessing the validity of a numerical calculation, if the analytic constant is known:

THEOREM 2.1. *Let a well-posed linear problem be given in the sense of* (1.2). *Any numerical technique that produces approximate solutions with small residuals will automatically guarantee small errors in the solution.* □

Though this theorem is trivial, it is important for providing a safe a-posteriori foundation for many unsafe and ad-hoc numerical calculations published in science and engineering journals. If the underlying problem is continuously dependent on the data and if the naive user at least checks the residuals, the calculations are on the safe side. But, unfortunately, there is no handbook listing all known inequalities of the form (1.2) for typical applications in science and engineering. In particular, it would be very useful to have proven upper bounds for the analytic constants.

But we also want to provide a special example for the rest of the paper. Guided by regularity theory for elliptic problems, we specialize to problems where the map L splits into maps $L^1, \ldots, L^k$ with $L^j : U \to F^j$, $1 \le j \le k$ such that $F = F^1 \times \ldots \times F^k$ is the data space. We assume $U = W_2^\mu(\Omega)$ for some bounded domain $\Omega \subset \mathbb{R}^d$ and $F^j := W_2^{\mu - f_j}(\Omega^j)$ where f_j is defined via a trace theorem by the order of the operator L^j and the dimension of the partial domain Ω^j. The space F is then equipped with the sup of the norms of the spaces F^j. The regularity subspace U_R occurring later will then be a space $U_R \subseteq W_2^m(\Omega) \subset U := W_2^\mu(\Omega)$ for some $m \ge \mu$.

In the standard Poisson problem with Dirichlet data we can take $L(u) = (-\Delta u, u_{\partial\Omega})$ mapping $U = W_2^\mu(\Omega)$ onto $F = W_2^{\mu-2}(\Omega) \times W_2^{\mu-1/2}(\partial\Omega)$. This is a well-established continuous dependence setting, if the domain is smooth enough to make trace theorems and the regularity order μ valid. See e.g. [15, 3] for early references which also allow distributional data and negative μ.

3. Approximation From Trial Spaces. The second ingredient of our framework is some additional regularity defined via a subspace U_R of the space U occurring in the continuous dependence bound (1.2). At this point, the regularity space can be quite general, but we also want an approximation property like (1.3) to hold. Thus we now have to consider our third ingredient, i.e. techniques that construct approximate solutions $\tilde{u}_r$ from a scale of *trial spaces* $U_r \subseteq U$ with a trial discretization parameter r. Note that this still includes plenty of methods, with or without meshes, like finite elements, Petrov-Galerkin schemes, spectral methods, and all variations of collocation. It is trivial that the choice of the trial space should be such that the true solution u can be approximated easily by functions from the trial space. In case of solutions with singularities, like for Poisson problems on domains with incoming vertices, one should make sure that the trial space contains the expected singular functions.

One way to make this more precise is to assume that there is a mapping $A_r : U_R \to U_r$ with

$$\|u - A_r(u)\|_U =: \delta_r(u) \text{ for all } u \in U_R \tag{3.1}$$

with a certain error $\delta_r(u)$ which will depend on the regularity subspace U_R.

But the previous section teaches us that we do not need to approximate the exact solution u in the space U by functions $u_r \in U_r \subset U$ directly. It suffices to make sure that the residuals $L(u) - L(u_r)$ are small. Thus the crucial quantity is the residual error $\|L(u - u_r)\|_F$ for any $u \in U$ and an approximation $u_r \in U_r$. In contrast to the theory of finite elements, we do not consider optimal approximations of u by u_r here, nor do we attempt to minimize the above error with respect to u_r. Instead, we are satisfied if the trial space U_r contains for each function $u \in U$ an approximation $u_r := I_r(u)$ with small *residual error* $\epsilon_r(u)$ as in (1.3).

Of course, if an L-independent approximation operator A_r with (3.1) is available, one can take $I_r = A_r$ and assume $\epsilon_r(u) \le \|L\|\delta_r(u)$ because of

$$\epsilon_r(u) = \|L(u - I_r(u))\|_W \le \|L\|\|u - A_r(u))\|_U \le \|L\|\delta_r(u).$$

But there may be better choices of I_r if L and the special structure of the residual space W are taken into account.

Inspection of (1.3) for F being a Cartesian product of Sobolev spaces reveals that the special approximation $I_r(u)$ should approximate u well *including its derivatives*, as far as they occur in the collection of data spaces F^j forming the space F. In fact, if we go back to our special case $U = W_2^\mu(\Omega)$ and $F = F^1 \times \ldots \times F^k$ with

$F^j := L^j(U) := W_2^{\mu-f_j}(\Omega^j)$ with a regularity subspace $U_R \subseteq W_2^m(\Omega)$ and $m > \mu$, we should expect approximation bounds like

$$\begin{aligned} \|L^j(u) - L^j(I_r(u))\|_{W_2^{\mu-f_j}(\Omega^j)} &\leq Cr^{m-\mu}\|L^j u\|_{W_2^{m-f_j}(\Omega^j)} \\ &\leq Cr^{m-\mu}\|u\|_{W_2^m(\Omega)} \text{ for all } u \in W_2^m(\Omega), \end{aligned} \tag{3.2}$$

and this should work for a reasonable choice of $0 \leq \mu \leq m$ and with rates that just depend on the regularity gap $m - \mu$, not on the order of the operators involved.

Note that the standard trial spaces of h–type finite element techniques consisting of piecewise linear functions fail to provide approximations of more that first-order derivatives. In contrast to this, trial spaces generated by sufficiently smooth kernel functions can contain approximations to derivatives of any order, without any additional work needed. We shall explain this in the following sections.

In contrast to many engineering applications where a rather simple solution function is calculated via a huge FEM system of millions of unknowns, we tend to argue in favor of small trial spaces designed to capture the essential features of the solution without taking the detour via a fine-grained space discretization. The consequence will be that the linear systems get unsymmetric, because any solution from a small *trial* space must be *tested* on a fine-grained space discretization, asking for many more degrees of freedom on the "test side" than on the "trial side". Unsymmetry of a method can be a feature instead of a bug. In what follows we shall investigate the relation of test and trial spaces more closely.

4. Kernel-Based Trial Spaces. Now it is time to study maps I_r or A_r with good approximation properties for certain trial spaces U_r in the sense of (1.3) and (3.1). This is independent of PDE solving, and we shall see that nonstationary meshless kernel-based trial spaces work perfectly.

DEFINITION 4.1. *A* kernel *is a function of the form* $K : \Omega \times \Omega \to \mathbb{R}$ *with* $\Omega \subseteq \mathbb{R}^d$. *It is* translation-invariant, *if* $K(x,y) = \Phi(x-y)$ *with* $\Omega = \mathbb{R}^d$ *and* $\Phi : \mathbb{R}^d \to \mathbb{R}$. *It is* radial *if it is translation-invariant and of the form*

$$K(x,y) = \Phi(x-y) = \phi(\|x-y\|_2) \text{ with } \phi : [0,\infty) \to \mathbb{R} \text{ and } x,y \in \mathbb{R}^d.$$

Radial kernels are also called radial basis functions.

Note that radial basis functions ϕ can in principle be used in any space dimension, but certain properties of the associated translation–invariant kernel Φ on $\mathbb{R}^d$ may depend on the dimension d.

Kernels provide excellent tools in various disciplines, including Approximation Theory, Partial Differential Equations, and Machine Learning. The most important kernels are *reproducing kernels* of some Hilbert space which can be called the "native" Hilbert space for the kernel. Any Hilbert space H of functions on a domain Ω with continuous and linearly independent point evaluations has a kernel K with the reproduction property

$$f(x) = (f, K(x,\cdot)_r \text{ for all } f \in H,\ x \in \Omega.$$

Conversely, any (strictly) positive definite [6, 22] and continuous kernel K on Ω is the reproducing kernel of a *native* Hilbert space N_K of continuous functions on Ω. We denote the norm on the native space N_K by $\|.\|_K$.

We focus here on *trial spaces* provided by kernels. Like in wavelet theory, the notions of *translation* and *dilation* play an important role. First, a general kernel

$K : \Omega \times \Omega \to \mathbb{R}$ can be translated to points $y_1, \ldots, y_M \in \Omega$ called *centers* to provide trial functions $u_j(x) := K(x, y_j),\ 1 \le j \le M$ on Ω.

In many cases, the set $Y := \{y_j,\ 1 \le j \le M\}$ of centers should fill a bounded domain Ω in such a way that the centers get dense when $M \to \infty$. This is expressed by the *fill distance*

$$h := h(Y, \Omega) := \sup_{x \in \Omega} \inf_{y \in Y} \|x - y\|_2$$

depending on Ω and the M centers in Y, and which should converge to zero if M tends to infinity. The fill distance is the radius of the largest open ball in Ω that contains none of the centers y_j from Y. We use the notation h here, but later we shall have two different fill distances for trial and test centers, and then we shall use r and s for clarity.

A *nonstationary* kernel-based trial space can now be defined as

$$U_r := \text{span}\ \{K(\cdot, y_j) : 1 \le j \le M\} \text{ with } r := h := h(Y, \Omega) \tag{4.1}$$

where the dependence on the location and number of the centers is suppressed in the notation.

In the above *nonstationary* situation, only translations, but no dilations are used. The translated kernel is fixed and independent of the fill distance $h = r$. There is no rescaling, if $h = r$ gets small. This is in contrast to the *stationary* technique in standard and general finite elements [4]. There, the basis functions are rescaled when the fill distance changes, and in the translation–invariant kernel–based case this can be described by

$$U_r := \text{span}\ \left\{ \Phi\left(\frac{x - y_j}{r} \right) : 1 \le j \le M \right\}$$

where now the wavelet-style interaction of translation and dilation are apparent.

The mathematics of the stationary and nonstationary case are quite different. This often leads to misunderstandings. The stationary situation, as included in the Generalized Finite Element Method [4], uses polynomial reproduction and the Bramble–Hilbert Lemma. If centers are on a grid, it applies the Strang–Fix theory. Convergence orders are closely tied to polynomial reproduction properties, and the choice of kernels is quite restricted, because integrable kernels like the Gaussian do not yield convergent stationary approximations for $h \to 0$ [6]. Stability is usually much better than in the nonstationary case, but convergence rates are much smaller. We focus on nonstationary kernel–based trial spaces here, because condition problems can be overcome [5, 12], and we are heading for methods with high approximation orders.

We define a map $I_r : u \to u_r := I_r(u) \in U_r$ of (1.3) via interpolation in the centers by solving the system

$$u_r(y_k) := \sum_{j=1}^{M} \alpha_j K(y_j, y_k) = u(y_k),\ 1 \le k \le M \tag{4.2}$$

for the coefficients $\alpha_1, \ldots, \alpha_M$ defining the function $u_r := I_r(u)$ in terms of the basis functions of the nonstationary trial space U_r of (4.1). This interpolation problem is solvable by definition, if the kernel K is symmetric and positive definite [6, 22],

Function	$\phi(r)$	Range	Smoothness β
Gaussian	$\exp(-r^2)$	$d \geq 1$	all β
inverse multiquadric	$(r^2+c^2)^\gamma,\ \gamma < -d/2,\ c > 0$	$d \geq 1$	all β
Sobolev for $W_2^k(\mathbb{R}^d)$	$r^{k-d/2}K_{k-d/2}(r),\ k > d/2$	$d \geq 1$	$\beta = 2k-d$
Wendland C^2 [20]	$(1-r)_+^4(1+4r)$	$d \leq 3$	3
	$(1-r)_+^5(1+5r)$	$d \leq 5$	3
Wendland C^4 [20]	$(1-r)_+^6(3+18r+35r^2)$	$d \leq 3$	5
	$(1-r)_+^7(1+7r+16r^2)$	$d \leq 5$	5

TABLE 4.1
Radial basis functions $\phi(r)$, positive definite on $\mathbb{R}^d$

because then the $M \times M$ matrix with entries $K(y_j, y_k)$ is symmetric and positive definite. Table 4.1 gives some examples. We ignore *conditionally* positive definite kernels here and refer to [6, 18, 22] for details.

The interpolation system (4.2) makes sense for all functions u which have well-defined function values at the centers y_k. Thus the mapping I_r is at least defined on $C(\Omega)$, but for solutions u of PDE problems in strong form we use it on a regularity subspace U_R of $C(\Omega) \subset U$.

The book [22] contains a fairly complete account of interpolation error bounds in the nonstationary setting, while bounds for stationary and regular cases are dominant in [6]. But in view of (1.3) and (3.2), we need very general error bounds in Sobolev spaces which are not covered in these books. Here (on the trial side) and later (on the test side) we use a general result from [17] which was extended in [23].

THEOREM 4.2. *Fix a bounded domain $\Omega \subset \mathbb{R}^d$ with a Lipschitz boundary and an interior cone condition. Assume $u \in W_2^m(\Omega)$ and fix a constant μ with $0 \leq \mu < \lfloor m \rfloor - d/2$. Then there are constants $C > 0$ and $h_0 > 0$ such that*

$$|u|_{W_2^\mu(\Omega)} \leq C \left(h^{m-\mu} |u|_{W_2^m(\Omega)} + h^{-\mu} \|u\|_{\infty, Y_h} \right) \tag{4.3}$$

holds for all $u \in W_2^m(\Omega)$, all $0 < h \leq h_0$ and all finite subsets Y_h in Ω with fill distance at most h.

This can be seen as a quantitative Poincarè-Friedrichs inequality for functions which are small on a finite subset, and it is independent of any trial space.

But if we now take $h = r$ and interpolate a function u on Y_h by $I_r(u)$ using points from Y_h as translations in (4.1) we get

$$|u - I_r(u)|_{W_2^\mu(\Omega)} \leq C r^{m-\mu} |u - I_r(u)|_{W_2^m(\Omega)} \text{ for all } u \in W_2^m(\Omega). \tag{4.4}$$

Then we use the standard fact that the interpolant $I_r(u)$ solves the minimization problem

$$\|v\|_K \to \min,\ v \in K,\ v(y_j) = u(y_j) \text{ for all } y_j \in Y_h,$$

implying that $\|I_r(u)\|_K \leq \|u\|_K$ holds if we assume u to be in the native space N_K for the kernel K.

Therefore we strengthen the requirement on the regularization subspace U_R and on the regularity of our solution u to

$$u \in N_K = U_R \subseteq W_2^m(\Omega) \subseteq U. \tag{4.5}$$

with bounded embeddings. This is easy if the kernel is smooth enough, and for the kernels in Table 4.1 the inequality $\beta + d \geq 2m$ is a sufficient condition.

We can now replace (4.4) by

$$|u - I_r(u)|_{W_2^\mu(\Omega)} \le Cr^{m-\mu}\|u\|_K \text{ for all } u \in N_K$$

with a different constant. This inequality can be coupled with trace theorems for the operators $L^j : W_2^m(\Omega) \to W_2^{m-f_j}(\Omega^j)$ to get

$$\|L^j(u - I_r(u))\|_{W_2^{\mu-f_j}(\Omega^j)} \le C\|u - I_r(u)\|_{W_2^\mu(\Omega)} \le Cr^{m-\mu}\|u\|_K \text{ for all } u \in N_K$$

which yields (3.2) in a slightly restricted form and our third ingredient (1.3) as

$$\|L(u - I_r(u))\|_F \le Cr^{m-\mu}\|u\|_K =: \epsilon_r(u) \text{ for all } u \in U_R = N_K. \tag{4.6}$$

5. Stability of Kernel-Based Test Discretizations. We now consider the stability conditions (1.4), our fourth ingredient. This requires several steps, and we give a compressed form of the argument first, using the terminology of the introduction.

THEOREM 5.1. *Assume scales of trial subspaces $U_r \subseteq U$ and test spaces F_s with a-priori inequalities* (1.2) *and*

$$\begin{aligned} \|L(u)\|_F &\le C_1(s)\|u\|_{U_R} + C_2(s)\|\Pi_s L(u)\|_{F_s} \text{ for all } u \in U_R \subseteq U \\ \|u_r\|_{U_R} &\le C_3(r)\|u_r\|_U \text{ for all } u_r \in U_r \subseteq U \\ C_1(s)C_3(r)C_a &\le 1/2. \end{aligned} \tag{5.1}$$

Then, in view of (1.4)*, these scales satisfy*

$$\|L(u_r)\|_F \le 2C_2(s)\|\Pi_s L(u_r)\|_{F_s} \text{ for all } u_r \in U_r.$$

Proof. Just consider

$$\begin{aligned} \|L(u_r)\|_F &\le C_1(s)\|u_r\|_{U_R} + C_2(s)\|\Pi_s L(u_r)\|_{F_s} \\ &\le C_1(s)C_3(r)\|u_r\|_U + C_2(s)\|\Pi_s L(u_r)\|_{F_s} \\ &\le C_1(s)C_3(r)C_a\|L(u_r)\|_F + C_2(s)\|\Pi_s L(u_r)\|_{F_s}. \end{aligned}$$

□

To explain the meaning of this machinery, we note that the first a-priori bound of (5.1) is modeled after (4.3) of Theorem 4.2, but now on the test side with parameter s. The quantity $C_1(s)$ will be a positive power of s related to the smoothness difference between spaces F and U_R, while $C_2(s)$ can be a negative power of s depending on μ. The second is an inverse Bernstein-type inequality on the finite-dimensional trial space U_r, because the left-hand side norm is stronger than the right-hand side norm. Here, we must expect that $C_3(r)$ is a negative power of r. The third inequality means that the granularity of the test discretization must be fine enough in relation to the trial discretization in order to guarantee that the homogeneous linear system $\Pi_s L(u_r) = 0$ has full rank, as postulated by the assertion of the theorem. This is a quantitative version of an earlier density result in [13].

We now have to prove the a-priori inequalities of (1.4) and (5.1) one by one. We do this for meshless kernel-based trial spaces and for our running example generalizing the Poisson equation. The trial discretization via U_r and a set Y_r of centers is chosen as in §4. We assume (4.5) and have the approximation result (4.6). On the test side, we use a set X_s of test centers which has a fill distance s on all of $\overline{\Omega}$. For all the

operators L^j that arise in L, we will have a selection $X_s^j := X_s \cap \Omega^j$ of points with the same fill distance with respect to Ω^j, because we can assume that all Ω^j are subsets of $\overline{\Omega}$. The projectors Π_s^j on F^j just map functions from $F^j = W_2^{\mu-f_j}(\Omega^j)$ to their values on X_s^j. We thus have to assume Sobolev embedding conditions

$$2(\mu - f_j) > \dim(\Omega^j),\ 1 \le j \le k. \tag{5.2}$$

The discretized spaces F_s^j will be $\mathbb{R}^{|X_s^j|}$ with the L_∞ norm, and we have

$$\Pi_s L^j(u) = (L^j(u))(X_s^j),\ \|\Pi_s L(u)\|_{F_s^j} = \|L^j(u)\|_{\infty,X_s^j}.$$

This implies by Sobolev embedding

$$\|\Pi_s L(u)\|_{F_s^j} = \|L^j(u)\|_{\infty,X_s^j} \le \|L^j(u)\|_{\infty,\Omega^j} \le C\|L^j(u)\|_{W_2^{\mu-f_j}(\Omega^j)}$$

where the constant is independent of u and s. We now assemble this into a discretization $F_s := F_s^1 \times \cdots \times F_s^k$ with $\Pi_s := \Pi_s^1 \times \cdots \times \Pi_s^k$ of $F = F^1 \times \cdots \times F^k$ and take the sup norm of the components. Then we have

$$\begin{aligned}
\|\Pi_s L(u)\|_{F_s} &= \sup_{1\le j\le k} \|\Pi_s^j L^j(u)\|_{F_s^j} \\
&= \sup_{1\le j\le k} \|L^j(u)\|_{\infty,X_s^j} \\
&\le \sup_{1\le j\le k} \|L^j(u)\|_{\infty,\Omega^j} \\
&\le C \sup_{1\le j\le k} \|L^j(u)\|_{W_2^{\mu-f_j}(\Omega^j)} \\
&= C\|L(u)\|_F
\end{aligned}$$

and get the second inequality of (1.4) with a constant that is independent of s and only dependent on Sobolev embedding. This leaves us to prove (5.1) in order to get the first inequality of (1.4).

Fortunately, the inequality (4.3) holds for general Sobolev spaces, and we can apply it on the test side for different operators. In fact,

$$\begin{aligned}
\|L^j(u)\|_{W_2^{\mu-f_j}(\Omega^j)} &\le C\left(s^{m-\mu}\|L^j(u)\|_{W_2^{m-f_j}(\Omega^j)} + s^{-(\mu-f_j)}\|L^j(u)\|_{\infty,X_s^j}\right) \\
&\le C\left(s^{m-\mu}\|u\|_{W_2^m(\Omega)} + s^{-\mu}\|L^j(u)\|_{\infty,X_s^j}\right) \\
\|L(u)\|_F &\le C\left(s^{m-\mu}\|u\|_K + s^{-\mu}\|\Pi_s L(u)\|_{F_s}\right).
\end{aligned}$$

We now go for the inverse inequality in (5.1). Assume K to be a translation-invariant kernel of finite smoothness which is Fourier-transformable in $\mathbb{R}^d$ with an exact decay

$$c(1+\|\omega\|_2)^{-\beta-d} \le \hat{K}(\omega) \le C(1+\|\omega\|_2)^{-\beta-d} \text{ for all } \omega \in \mathbb{R}^d \tag{5.3}$$

where the constants β can be read off Table 4.1. Then we can cite the Bernstein-type inequality

$$\|u_r\|_{W_2^{(d+\beta)/2}(\mathbb{R}^d)} \le Cr^{-(d+\beta)/2}\|u_r\|_{L_2(\mathbb{R}^d)} \text{ for all } u_r \in U_r$$

from [16] provided that the trial centers in Y_r are not too wildly scattered in the sense that the minimal separation distance $q(Y_r)$ is uniformly bounded below by the fill distance $h(Y_r,\Omega)$ via

$$q(Y_r) := \min_{y_j \ne y_i \in Y_r} \|y_j - y_i\|_2 \ge C \sup_{y\in\Omega} \min_{y_i\in Y_r} \|y - y_i\|_2 =: h(Y_r,\Omega) \tag{5.4}$$

such that both quantities behave asymptotically like the trial discretization parameter r. The follow-up paper [19] extends this to

$$\|u_r\|_{W_2^m(\Omega)} \le Cr^{\mu-m}\|u_r\|_{W_2^\mu(\Omega)} \text{ for all } u_r \in U_r \tag{5.5}$$

with the range $0 \le \mu \le m \le \beta$, and this is precisely what we need.

6. Strong Convergence in Sobolev Spaces. We now assemble what we have in case of our running example with continuous dependence in Sobolev norms. In contrast to the introduction, we proceed here from the user's point of view.

We start with the analytic problem. Consider an operator equation $L(u) = f$ as in (1.1) whose solution u is continuously dependent on the data f. The analytic constant C_a of (1.2) is assumed to hold if we pick spaces $U = W_2^\mu(\Omega)$ and F defined as a Cartesian product of Sobolev trace spaces as in §2. If several choices of μ are possible, the user should know that the final convergence will take place in $U = W_2^\mu(\Omega)$, but large μ have to be paid for by regularity. If convergence of higher-order derivatives is of importance, a sufficiently large μ should be chosen. Since we solve problems in strong form via evaluation of residuals, we have to pick μ large enough to let all data have continuous point evaluations. This is expressed by the requirement (5.2). At this point, the lower bounds for μ will rule out problems with low regularity. Such problems should be tackled with methods using weak data functionals and involving integration. We plan to deal with such methods in the future, in particular with the unsymmetric meshless Petrov-Galerkin method of Atluri and his collaborators [1].

The next step concerns regularity. We assume that the solution should have at least a $U_R := W_2^m(\Omega)$ regularity with some $m > \mu$. By standard arguments from Approximation Theory, the difference between m and μ is the driving force for the possible convergence rates. The user has to decide which m is adequate. Larger m will improve the convergence rates, but they may not be justified by the smoothness of the problem.

Then we pick a kernel K which is smooth enough to have its native space N_K contained in $W_2^m(\Omega)$. In view of Table 4.1, this requires $\beta + d \geq 2m$. The solution u must have some more regularity than $W_2^m(\Omega)$, because it should be in $U_R = N_K$. The excess regularity of N_K over $W_2^m(\Omega)$ does not pay off later, and thus it is a good idea to stay with a kernel satisfying $\beta + d = 2m$ to have norm equivalence between $W_2^m(\Omega)$ and N_K. Note that the compactly supported radial polynomial kernels of Wendland [20] satisfy this for certain choices of m, β, and d.

Now it is time to pick a meshless trial discretization U_r via a set Y_r of trial centers with fill distance r using the kernel K. Then we could expect the convergence rate $\epsilon_r(u) \leq Cr^{m-\mu}$ for $u \in N_K$ if we could directly interpolate the solution in the points of Y_r. The constants $C_3(r)$ in (5.1) will then come out to be of order $r^{\mu-m}$ via (5.5). However, for this we have to assume (5.3) and (5.4), meaning that the kernel should have finite smoothness and that the trial centers are not too wildly scattered.

The next step is to pick the test discretization via a set X_s of test centers with fill distance s. The first inequality of (5.1) will then hold with $C_1(s) \leq Cs^{m-\mu}$ and $C_2(s) \leq Cs^{-\mu}$. The second inequality of (1.4) holds with $c(s)$ independent of s because we assume continuous residuals and corresponding Sobolev embedding theorems. But we have to make the test discretization fine enough to satisfy the third inequality of (5.1), which amounts to

$$Cs^{m-\mu}r^{\mu-m} < \frac{1}{2} \text{ or } s < r \cdot \frac{1}{(2C)^{m-\mu}}. \tag{6.1}$$

As expected, this means that the test discretizations must be somewhat finer than the trial discretizations, but the ratio of the granularities of the discretizations can be kept fixed. There is plenty of leeway for small trial and large test spaces.

We are now ready to put everything into Theorem 1.1, while we assume that we solve the discretized problem (1.5) with accuracy $\delta_{r,s}$. With new generic constants

we get

$$\|u - u^*_{r,s}\|_{W_2^\mu(\Omega)} \quad \leq \quad C\,(r^{m-\mu}\,(1 + Cs^{-\mu}) + \delta_{r,s})\,\|u\|_K. \tag{6.2}$$

Note that this bound has the proper approximation error of order $r^{m-\mu}$ holding between Sobolev spaces $U = W_2^\mu(\Omega)$ and $U_R = N_K \subseteq W_2^\mu(\Omega)$, but there also is a counteracting term $s^{-\mu}$ which is the price we have to pay for working on discrete residuals in the L_∞ norm while bounding the error in the norm on $U = W_2^\mu(\Omega)$. If we choose $\delta_{r,s}$ properly via (1.6) and s via (6.1), we have solvability of the system and an error bound

$$\|u - u^*_{r,s}\|_{W_2^\mu(\Omega)} \leq Cr^{m-2\mu}\|u\|_K.$$

7. Weak Convergence in Sobolev Spaces. Analysis of the previous section shows that the term $s^{-\mu}$ in the first line of (6.2) with some positive μ satisfying (5.2) makes the final bound worse than what could be expected as an optimal result. Tracing this back to (4.3) shows that one should better look at another variation which allows $\mu = 0$ without spoiling the assumption that the data are still continuous. In fact, reference [23] also has

$$|u|_{L_\infty(\Omega)} \leq C\left(h^{m-d/2}|u|_{W_2^m(\Omega)} + \|u\|_{\infty,Y_h}\right) \text{ for all } u \in W_2^m(\Omega) \tag{7.1}$$

if $\lfloor m \rfloor > d/2$. But this does not easily fit into the framework required for continuous dependence unless we make sure that all the spaces F^j carry the L_∞ norm and are still useful for continuous dependence. But this would mean that we take μ so small that we have a continuous embedding $C(\Omega^j) \subseteq W_2^{\mu-f_j}(\Omega^j)$, i.e. we now take $\mu = \min_j f_j$. Note that this will lead to a weak convergence result in $U := W_2^\mu(\Omega)$, though the problem formulation is still strong. For instance, a problem with Dirichlet data will lead to $\mu = 1/2$ due to the trace map $W_2^\mu(\Omega) \to W_2^{\mu-1/2}(\partial\Omega)$ if all other trace or differential operators have a larger loss in the order of the respective Sobolev trace spaces.

Thus we now repeat our basic argument for $U := W_2^\mu(\Omega)$ with $\mu = \min_j f_j$. By the standard embedding theorems we now choose the spaces F^j as $C(\Omega^j)$ with the L_∞ norm and still get continuous dependence, provided that the analytic problem allows this choice of μ. This is not trivial, because right-hand sides of differential equations now have to allow distributional data in general. But for the Poisson problem this is correct, if the domain is smooth [15, 3].

Our choice of regularity space U_R and the kernel K will be as above. This fixes β and m. Then we have to look at the approximation order in (1.3) and get

$$\|u - I_h(u)\|_{W_\infty^\ell(\Omega)} \leq Cr^{m-\ell-d/2}\|u\|_K \text{ for all } u \in U_R = N_K$$

and $0 \leq \ell < \lfloor m \rfloor - d/2$. If we combine this with trace theorems for the operators F^j we find

$$\|L_j(u - I_h(u))\|_{L_\infty(\Omega^j)} \leq C\|u - I_h(u)\|_{W_\infty^{f_j}(\Omega^j)} \leq Cr^{m-f_j-d/2}\|u\|_K \text{ for all } u \in U_R = N_K.$$

Thus we get

$$\epsilon_r(u) \leq Cr^{m-d/2-\max_j(f_j)}$$

for (1.3). The second inequality of (5.1) can handled as in the previous section, giving $C_3(r) \le Cr^{\mu-m}$.

The discretization of the F^j spaces is again by pointwise evaluation on a set X_s of test centers, taking the discrete L_∞ norm. Since the spaces F_s^j are now equipped with the L_∞ norm, we have $c(s) = 1$ in the second inequality of (1.4). The proof of the first inequality of (5.1) now starts with (4.3) on the various data:

$$\begin{aligned} |L_j(u)|_{L_\infty(\Omega^j)} &\le C\left(s^{m-f_j-d_j/2}|L_j(u)|_{W_2^{m-f_j}(\Omega^j)} + \|L_j(u)\|_{\infty,X_s}\right) \\ &\le C\left(s^{m-f_j-d_j/2}\|u\|_{W_2^m(\Omega)} + \|L_j(u)\|_{\infty,X_s}\right) \\ &\le C\left(s^{m-f_j-d_j/2}\|u\|_K + \|L_j(u)\|_{\infty,X_s}\right) \end{aligned}$$

for all $u \in U_R = N_K$ where d_j is the dimension of Ω^j. Thus the first inequality of (5.1) holds with $C_2(s)$ independent of s and with

$$C_1(s) \le Cs^{m-d/2-\max_j f_j}.$$

Thus we get uniform stability of the discretization scheme, if the third inequality

$$Cs^{m-d/2-\max_j f_j} r^{\mu-m} < \frac{1}{2}$$

of (5.1) is satisfied. This holds if

$$s < Cr^{\frac{m-\min_j f_j}{m-d/2-\max_j f_j}}$$

which is not too bad for large m.

Since the discretization scheme is uniformly stable, Theorem 1.1 now gives the error bound

$$\|u - u_{r,s}^*\|_{W_2^{\min_j f_j}(\Omega)} \le Cr^{m-d/2-\max_j f_j}\|u\|_K \text{ for all } u \in N_K \subseteq W_2^m(\Omega)$$

provided that the numerical solution of (1.5) observes (1.6). The left-hand norm is rather weak here, and the approximation order can probably be improved. For instance, a standard two-dimensional Poisson problem with Dirichlet data would lead to

$$\|u - u_{r,s}^*\|_{W_2^{1/2}(\Omega)} \le Cr^{m-3}\|u\|_K \text{ for all } u \in N_K \subseteq W_2^m(\Omega),$$

but an optimal rate for the Sobolev spaces involved would be $m-1/2$ instead of $m-3$.

8. Numerical Methods. We now want to look at techniques to solve the discrete problem (1.5). It amounts to solve the k linear problems

$$\Pi_s^j L^j(u - u_{r,s}^*) = 0,\ 1 \le j \le k$$

approximately, where we discretized the operators L^j on the domains Ω^j by taking only point evaluations. This takes the form of collocation

$$L^j(u)(x_{ji}) = L^j(u_{r,s}^*)(x_{ji}) = 0,\ 1 \le j \le k,\ 1 \le i \le N_j$$

where the points of the test discretization X_s are the union of the sets

$$X_s^j := \{x_{j1}, \ldots, x_{jN_j}\},\ 1 \le j \le k,$$

and where we dropped the dependence on s in the notation for the x_{ji} and for N_j. For a shorthand notation, we introduce the functionals

$$\lambda_{ji} \ : \ v \mapsto L^j(v)(x_{ji})$$

and rearrange them into a single-indexed list $\lambda_1, \ldots, \lambda_N$ with $N = N_1 + \ldots + N_k$.

Since $u^*_{r,s}$ should be in the trial space U_r generated by translations of the kernel K at trial centers forming $Y_r := \{y_1, \ldots, y_M\}$, we arrive at a system

$$\sum_{m=1}^{M} \alpha_m \lambda_i^z K(z, y_m) = \lambda_i^z u(z), \ 1 \le i \le N$$

with M unknowns and N equations. In case $M = N$ this is exactly the unsymmetric collocation technique dating back to E. Kansa in 1986 [11]. It has no rigid foundation yet, and it can fail in specially constructed situations [10], though is works fine in many applications. In the first years it was applied to small problems with smooth solutions due to serious condition problems, but recently there are results on preconditioning [5, 12] that allow a wider range of applicability.

In view of Theorem 5.1 and the two previous sections we know that $N \ge M$ holds and the system has full rank M. Thus it will be unsymmetric and overdetermined, but at least there is no rank loss. Furthermore, we know by (1.3) and (1.6) that there is a good approximate solution to the full system. This means that we can allow any numerical method that produces a solution with similar or less deviation.

Since our convergence analysis worked with the L_∞ norm on the discretized F_s spaces, a first choice would be to go for a best L_∞ approximation of the right-hand side. This means to solve a linear optimization problem which minimizes η under the constraints

$$-\eta \le \sum_{m=1}^{M} \alpha_m \lambda_i^z K(z, y_m) - \lambda_i^z u(z) \le \eta, \ 1 \le i \le N \tag{8.1}$$

where $\alpha_1, \ldots, \alpha_M$ are the other variables. If the revised simplex method is applied to the dual problem, each step has a $\mathcal{O}(M^2)$ complexity. The Kuhn-Tucker conditions ensure that one can work with at most $M+1$ active test conditions at each time. This makes the number $N >> M$ of test centers much less relevant than M, and for nicely chosen low-dimensional trial spaces one can get away with rather small computational complexity. We give an example in the final section.

But one can also try all other techniques that somehow provide a function $u^*_{r,s} \in U_r$ which by a-posteriori inspection leads to a small residual norm $\delta_{r,s}$ in (1.5). This can happen to the original Kansa method when executed on a subset of M test points, or by adaptive bootstrapping techniques like the one in [14, 13] which picks suitable test centers and trial centers one-by-one. Other alternatives are to use pivoting with row exchange or to go for a least-squares solution first. Anyway, if the resulting residual norm $\delta_{r,s}$ is small, the result of Theorem 1.1 is still valid, proving that one actually has a good approximation to the real solution. An illustration follows in the final section.

As an aside, we note that that a simpler theory is possible, if we optimize over a non-discrete residual norm on F. Section 5 will then be obsolete, but one has to solve semi-infinite optimization problems (if F carries a sup-norm) or apply least-squares methods with integrations (if F carries an inner product).

9. Ill-posed Problems. For ill-posed problems, continuous dependence fails, but our method and its analysis will still be useful. We assume that the problem still has the form (2.3), but we now assume that the "true solution" $u \in U$ only satisfies

$$L(u) = f + \rho \in F \tag{9.1}$$

where F contains the available data f and a small residual ρ. The problem $L(u) = f$ may be unsolvable. We consider a function $\tilde{u} \in U$ to be acceptable as a solution, if

$$\|L(\tilde{u}) - L(u)\|_F = \|L(\tilde{u}) - f - \rho\|_F$$

is not much larger than $\|\rho\|_F$. We still assume (1.3) and (1.4), but we have to replace (1.5)by

$$\|\Pi_s(f - L(u^*_{r,s}))\|_{F_s} \le \delta_{r,s}, \tag{9.2}$$

because $L(u)$ now is unknown and does not coincide with f. Furthermore, solvability of the above system now requires

$$c(s)(\|\rho\|_F + \epsilon_r(u)) \le \delta_{r,s} \tag{9.3}$$

instead of (1.6) as a sufficient condition. The proof technique of Theorem 1.1 then still implies

THEOREM 9.1. *If the analytic problem is ill-posed, but solvable by $u \in U_R$ in the sense of* (9.1)*, and if we solve* (9.2) *by some $u^*_{r,s} \in U_r$, then there is a bound*

$$\|L(u - u^*_{r,s})\|_U \le c(s)\|\rho\|_F + \left(\epsilon_r(u)\left(1 + \frac{C(r,s)}{c(s)}\right) + c(s)\delta_{r,s}\right).$$

If the discretization is uniformly stable, then there is a choice of $\delta_{r,s}$ via (9.3) *such that the above residual error behaves asymptotically like the trial approximation error $\epsilon_r(u)$ plus $\|\rho\|_F$.*

Proof. We modify the proof of Theorem 1.1 to get

$$\begin{aligned}
\|L(u - u^*_{r,s})\|_F &\le \|L(u - I_r(u))\|_F + \|L(I_r(u) - u^*_{r,s})\|_F \\
&\le \epsilon_r(u) + c(s)\|\Pi_s L(I_r(u) - u^*_{r,s})\|_{F_s} \\
&\le \epsilon_r(u) + c(s)\|\Pi_s L(I_r(u) - u)\|_{F_s} \\
&\quad + c(s)\|\Pi_s(L(u) - f)\|_{F_s} \\
&\quad + c(s)\|\Pi_s(f - L(u^*_{r,s}))\|_{F_s} \\
&\le \epsilon_r(u) + c(s)\delta_{r,s} + \tfrac{C(r,s)}{c(s)}\|L(I_r(u) - u)\|_F + c(s)\|\Pi_s\|\|\rho\|_F \\
&\le c(s)\|\rho\|_F + \epsilon_r(u)\left(1 + \tfrac{C(r,s)}{c(s)}\right) + c(s)\delta_{r,s} \qquad \square
\end{aligned}$$

□

For simplicity of the above presentation, we have replaced the second inequality of (1.4) by

$$c(s)\|\Pi_s g\| \le \|g\|_F \text{ for all } g \in F$$

which is no serious complication. However, we should comment on what happens with Theorem 5.1 if we have no analytic constant C_a for (5.1). We replace C_a in the third inequality by the constant $C_4(r)$ arising in a finite-dimensional version

$$\|u_r\|_U \le C_4(r)\|L(u_r)\|_F \text{ for all } u_r \in U_r$$

of (1.2). This is feasible due to norm equivalence, but we leave it to future research to derive upper bounds for $C_4(r)$.

10. Examples. We close with a few simple examples. These are not intended to support questions concerning efficiency or the range of applicability. They only compare the strategies described in the previous section, and they come out of a rather simple MATLAB program that solves Poisson problems (2.1) on various domains with various data.

In view of our preference for small trial spaces and smooth problems, we focus on cases with a smooth solution which we use to generate the data f_Ω and f_D in (2.1). Since meshless collocation techniques can handle oddly-shaped regions easily, we take the domain of Figure 10.1 with polar representation

$$r = 0.3(1.5 + \cos(5\varphi)),\ \varphi \in [0, 2\pi)$$

and we now explain the notation there and in the following tables. In all cases, the exact solution of the Dirichlet problem is provided by the fundamental solution of the Laplacian centered at $(2, 2)$, i.e. outside the domain, but not too far away.

We generate test centers for Dirichlet boundary conditions (*Test DBC* in the figures) by mapping equidistant parameters from $[0, 2\pi)$ to the boundary. Test centers for the Laplacian (*Test PDE* in the figures) are selected by taking all points of a global grid that fall inside the domain. This is accomplished easily because we focus on star-shaped domains represented in polar coordinates. We add the locations of the Dirichlet boundary test centers, because we want the differential equation to be satisfied up to the boundary. Thus boundary test centers occur twice, because they serve for testing both the Dirichlet data and the differential equation. In the figures, we distinguish the Dirichlet boundary test data by circles from the interior test data for the PDE. There are no geometric calculations at all, and the test points have no information on their neighbors. The algorithms would run in quite the same way if data were scattered.

The choice of trial centers depends on the algorithm. For symmetric collocation (called *S fixed* in the tables, see references in [24, 9, 8]), the test functionals determine the trial functions, and there is nothing to choose. We just take our already fixed set of test centers. For the standard unsymmetric collocation due to E. Kansa (*K fixed* in the tables) the standard choice of trial centers is to take all test centers, if possible. This method works on a square unsymmetric system and hopes for nonsingularity. But since we use the boundary test centers twice, as Dirichlet boundary value test centers and as Laplacian test points falling on the boundary, we have to add intermediate boundary centers to the set of trial centers in order to have a square linear system. This point assignment is shown in Figure 10.1.

For the unsymmetric versions of collocation, we do not force the algorithms to use large fixed sets of trial or test centers right from the start. Instead, we run the technique from [14, 13] to select test and trial centers adaptively from a large set of choices described below. In short, the algorithm works on square systems of increasing size. To proceed to a larger system,

1. add a test functional where the current residual is largest in absolute value,
2. add a trial center that makes the determinant of the new system closest to one.

If there are $N >> n$ test functionals or $N >> n$ trial centers offered when going over from an $n \times n$ system to a $(n+1) \times (n+1)$ system, then the update complexity is $\mathcal{O}(n^2 \cdot N)$ in computation and storage. We leave details to [14, 13].

The choice of trial and test centers is made from a large and fine grid of centers that extends to $[-4, 4]^2$. Note that this range contains the singularity of the true

solution at the point $(2,2)$, but only the trial centers can be chosen close to the singularity. The test centers for the Laplacian are those that fall into $\overline{\Omega}$. In addition, we make sure that the test and trial centers of the methods *S fixed* and *K fixed* based on coarse fixed center selections are included. Larger and finer sets for the adaptive choice of centers do not change our results significantly, for reasons to become clear soon.

The greedy adaptive method of [13] is called *K Adap* in the tables, and it selects trial centers as well as test functionals from the large set described above. After it stops, we freeze the selected trial space and perform a Chebyshev-style minimization of the residuals on all possible test locations. This is method *K Opt* in the tables, and it is the linear optimization described as (8.1) in the previous section. For both adaptive methods, we weight the Dirichlet boundary data by a factor of 1000 in order to have some preference for the boundary test data over the interior test data, when it comes to selecting new test centers.

But now we have to explain our stopping criterion, which is somewhat unusual. Note that linear equation systems arising from nonstationary collocation problems have three typical properties:

1. the condition blows up when test or trial centers get close,
2. this blow-up grows with the smoothness of the kernel,
3. but the systems usually have a smaller subsystem that already has a useful solution. This is in sharp contrast to symmetric stiffness matrices arising from fine-grained spatial resolutions.

Thus one should always wisely use row and column selection strategies (e.g. pivoting or adaptive techniques), and stop the calculation at the condition limit, in order to finish with a solution of a "good" subsystem. Of course, this only holds if no other precautions like preconditioning [5, 12] are taken.

For all algorithms, we monitor the condition of the square (sub-) matrices that arise during the calculations. Algorithms are stopped (including pivoting) when the condition exceeds 10^{12}. We ignore preconditioning, because we want the condition problems of the various algorithms to show up correctly for comparison. Since we use condition as a stopping criterion, and since we know that there is an approximate solution with very small residuals, approximations with large residuals indicate premature stopping due to condition problems. Furthermore, refinement of grids for a very fine-grained selection of trial centers does not help much if there is already a good supply of trial centers on a more coarse grid that suffices to reach the condition limit. This is a reason on the trial side to show that it does not pay off significantly to offer much larger choices of centers. On the test side and for the optimizing technique, larger choices also do not make sense, because Kuhn-Tucker conditions will pick $\mathcal{O}(M)$ active test centers where the residual error alternates, and if these are refined locally on a very fine grid, there will be no reasonable decrease of the error.

It is true that our current theory requires trial centers to be restricted to the domain, while our examples allow trial centers outside the domain. Experience shows that restricted trial centers lead to a similar behavior, but to larger condition and earlier stopping. Since we found our adaptive algorithm to be wiser than the theory, we allowed trial centers outside the domain, being aware of the fact that this should be handled by the theory of follow-up papers.

In the tables we denote the final degrees of freedom in the approximate solution by d_u, while the number of offered trial and test points are d_{trial} and d_{test}. Premature stopping due to condition will lead to $d_u < \min(d_{trial}, d_{test})$ in the symmetric or in

the standard Kansa case. In all cases the maximum storage is of order $d_u(d_u + d_{trial} + d_{test})$, while computational complexity is of order $d_u^2(d_u + d_{trial} + d_{test})$, if primitive solvers are used. Note that the adaptive technique works "on-the-fly", while linear optimization uses the revised simplex method. A matrix of size $d_{trial} \cdot d_{test}$ is never needed, and there are no numerical integrations at all. But efficiency is not our major concern here. Large real-world problems should be split into smaller problems of this type by either partition-of-unity methods [21] or domain decomposition techniques [7, 12], and preconditioning along the lines of [5, 12] should be used together with iterative solvers.

Final error evaluation is done on much finer grids, and the results are called rd_∞ and rb_∞ for the sup norm error of residuals in the domain and on the boundary. By u_∞ we denote the evaluated sup norm error of the approximate solution. Note that if residuals are evaluated reliably, and if the analytic constant is known, the error bound (2.4) will apply without further work.

Table 10.1 shows results on the five-pointed domain using multiquadrics $\phi(r) = \sqrt{1.25^2 + r^2}$. As in the other tables, the differences in error should not be over-interpreted. They are strongly influenced by condition, because unstable methods must stop earlier, but of course there are also differences in the trial spaces. Note that the final solution for the adaptive and optimizing algorithms uses only 61 degrees of freedom, even if 499 test functionals and 8724 trial centers are possible. This means that there is a 61 by 61 subsystem of a full 499 by 8724 system such that the solution of the subsystem solves the full system with reasonable accuracy. Linear optimization still improves the quality of the other two non-square variations of Kansa's method by one decimal digit, while they share the same choice of 8724 trial centers to pick from. If plotted, the residuals after linear optimization show the standard equioscillation pattern known from Chebyshev approximation, while in the other cases the error is not well-distributed over the domain and the boundary.

Method	rd_∞	rb_∞	u_∞	d_u	d_{trial}	d_{test}
S fixed	2.55-4	2.43-6	1.22-6	72	75	75
K fixed	1.85-3	4.98-5	2.10-5	70	75	75
K Adap	2.01-5	1.13-7	1.07-7	61	8724	499
K Opt	3.62-6	1.27-8	1.06-8	61	8724	499

TABLE 10.1
Results for multiquadrics

For this example, Figure 10.1 shows the distribution of the test and trial centers for the standard Kansa method, while Figure 10.2 displays part of the trial and test centers picked by an adaptive choice of the trial space. Observe the dense distribution of test centers in $\overline{\Omega}$, while many of the adaptively chosen trial centers are outside of the domain.

Table 10.2 shows results on the five-pointed domain using Gaussians $\phi(r) = \exp(-r^2/1.25^2)$. The results are very similar to those of the multiquadrics.

11. Conclusions. We provided convergence proofs for a generalized non-square version of Kansa's collocation method, showing that the convergence rates are determined by approximation results for nonstationary meshless-kernel-based trial spaces. The rates improve with the smoothness of the solution, the domain, the differential operator, and the kernel. They hold for large classes of analytic problems, provided

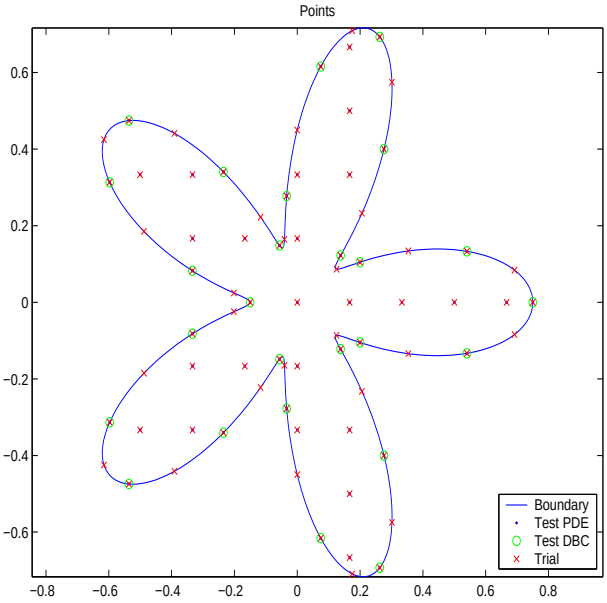

FIG. 10.1. *Fixed centers for Kansa's method*

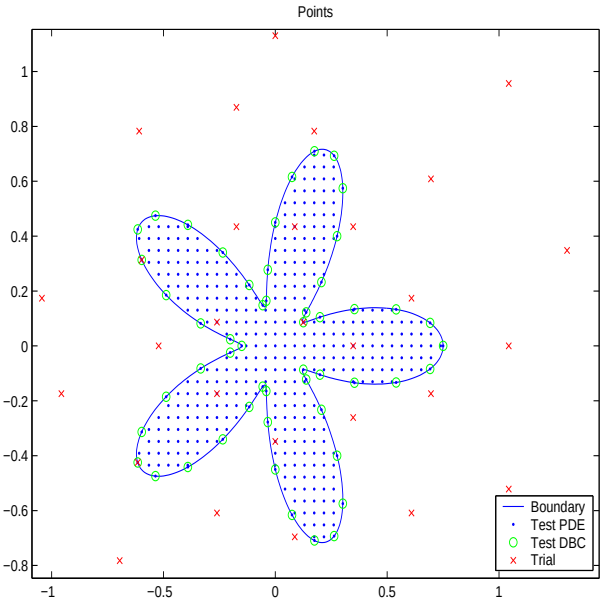

FIG. 10.2. *Centers after greedy optimization*

that there is continuous dependence on the data. On the downside, the naive practical application of nonstationary kernel-based approximation always faces condition problems, but new results demonstrate how to cope with these.

There are many possibilities for improvement and extension of these results:

1. Find sufficient conditions for nonsingularity of square Kansa-type collocation

Method	rd_∞	rb_∞	u_∞	d_u	d_{trial}	d_{test}
S fixed	1.90-5	2.06-6	9.63-7	62	75	75
K fixed	4.34-5	8.87-6	4.04-6	57	75	75
K Adap	7.28-6	8.11-8	6.79-8	60	8724	499
K Opt	1.72-6	5.20-9	7.45-9	60	8724	499

TABLE 10.2
Results for Gaussians

matrices.

2. Introduce discretization-dependent weights for different parts of residuals into the theory of this paper in order to align dimension- and order-dependent convergence rates.
3. For important problems of Applied Analysis, state the continuous dependence of the solution on the data in precise form and then derive the possible convergence rates for variations of Kansa's method, using the general theory of this paper.
4. Implement algorithms of this paper as local components of a global algorithm using localization features like domain decomposition or partitions of unity and efficiency-enhancing features like preconditioning and iterative solvers.
5. For such a global algorithm, perform large-scale numerical experiments and compare observed convergence rates with the theoretical ones of this paper.
6. Generalize all of this to unsymmetric methods for weak problems like the MLPG method of Atluri and collaborators [2].

Acknowledgement. The author thanks I. Babuška, C.S. Chen, Y.C. Hon, E. Kansa, L. Ling, G. Lube, and H. Wendland for several very stimulating discussions.

REFERENCES

[1] S. N. ATLURI AND T. L. ZHU, *A new meshless local Petrov-Galerkin (MLPG) approach in computational mechanics*, Computational Mechanics, 22 (1998), pp. 117–127.

[2] ———, *A new meshless local Petrov-Galerkin (MLPG) approach to nonlinear problems in computer modeling and simulation*, Computer Modeling and Simulation in Engineering, 3 (1998), pp. 187–196.

[3] J. AUBIN, *Approximation of elliptic boundary-value problems*, Pure and Applied Amathematics, vol. XXVI, Wiley-Interscience, 1972.

[4] I. BABUŠKA, U. BANERJEE, AND J. OSBORN, *Survey of meshless and generalized finite element methods a unified approach*, in Acta Numerica, 2003, pp. 1–215.

[5] D. BROWN, L. LING, E. KANSA, AND J. LEVESLEY, *On approximate cardinal preconditioning methods for solving pdes with radial basis functions*. Preprint, 2005.

[6] M. BUHMANN, *Radial Basis Functions*, Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, 2004.

[7] M. DUBAL, *Domain decomposition and local refinement for multiquadric approximations. I: Second-order equations in one-dimension*, J. Appl. Sci. Comp., 1 (1994), pp. 146–171.

[8] C. FRANKE AND R. SCHABACK, *Convergence order estimates of meshless collocation methods using radial basis functions*, Advances in Computational Mathematics, 8 (1998), pp. 381–399.

[9] ———, *Solving partial differential equations by collocation using radial basis functions*, Appl. Math. Comp., 93 (1998), pp. 73–82.

[10] Y. C. HON AND R. SCHABACK, *On unsymmetric collocation by radial basis functions*, J. Appl. Math. Comp., 119 (2001), pp. 177–186.

[11] E. J. KANSA, *Application of Hardy's multiquadric interpolation to hydrodynamics*, in Proc. 1986 Simul. Conf., Vol. 4, 1986, pp. 111–117.

[12] L. Ling and E. Kansa, *Preconditioning for radial basis functions with domain decomposition methods.* Preprint, 2005.

[13] L. Ling, R. Opfer, and R. Schaback, *Results on meshless collocation techniques.* Preprint on
`http://www.num.math.uni-goettingen.de/schaback/research/group.html`, 2004.

[14] L. Ling and R. Schaback, *On adaptive unsymmetric meshless collocation*, in ICCES04 proceedings, 2004.

[15] J. Lions and M. E., *Problèmes aux limites non homogènes at applications vol. 1*, Travaux et recherches mathématiques, Dunod, 1968.

[16] F. Narcowich, J. Ward, and H. Wendland, *Sobolev error estimates and a Bernstein inequality for scattered data interpolation via radial basis functions.* Preprint, 2005.

[17] F. J. Narcowich, J. D. Ward, and H. Wendland, *Sobolev bounds on functions with scattered zeros, with applications to radial basis function surface fitting*, Mathematics of Computation, (2003).

[18] R. Schaback and H. Wendland, *Characterization and construction of radial basis functions*, in Multivariate Approximation and Applications, N. Dyn, D. Leviatan, D.and Levin, and A. Pinkus, eds., Cambridge University Press, 2001, pp. 1–24.

[19] R. Schaback and H. Wendland, *Local Bernstein inequalities for radial basis functions.* Preprint, 2005.

[20] H. Wendland, *Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree*, Advances in Computational Mathematics, 4 (1995), pp. 389–396.

[21] ———, *Fast evaluation of radial basis functions: Methods based on partition of unity*, in Approximation Theory X: Wavelets, Splines, and Applications, C. K. Chui, L. L. Schumaker, and J. Stöckler, eds., Vanderbilt University Press, 2002, pp. 473–483.

[22] ———, *Scattered Data Approximation*, Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, 2005.

[23] H. Wendland and C. Rieger, *Approximate interpolation.* Preprint Göttingen, 2005.

[24] Z. Wu, *Hermite–Birkhoff interpolation of scattered data by radial basis functions*, Approximation Theory and its Applications, 8/2 (1992), pp. 1–10.

Institut für Numerische und Angewandte Mathematik
Universität Göttingen
Lotzestr. 16-18
D - 37083 Göttingen

Telefon: 0551/394512
Telefax: 0551/393944

Email: trapp@math.uni-goettingen.de URL: http://www.num.math.uni-goettingen.de

Verzeichnis der erschienenen Preprints:

2005-01	A. Schöbel, S. Scholl:	Line Planning with Minimal Traveling Time
2005-02	A. Schöbel	Integer programming approaches for solving the delay management problem
2005-03	A. Schöbel	Set covering problems with consecutive ones property
2005-04	S. Mecke, A. Schöbel, D. Wagner	Station location - Complex issues
2005-05	R. Schaback	Convergence Analysis of Methods for Solving General Equations
2005-06	M. Bozzini, L. Lenarduzzi, R. Schaback	Kernel B–Splines and Interpolation
2005-07	R. Schaback	Convergence of Unsymmetric Kernel–Based Meshless Collocation Methods