Compactly Supported RBF Interpolation vs. Spline Interpolation

Davoud Mirzaei

Department of IT, Uppsala University

March 16, 2025

1 Introduction

Compactly supported radial basis functions (RBFs) and spline basis functions may look very similar in one dimension: both are local, their supports may shrink proportionally to the mesh width, and in some special cases both satisfy a Kronecker delta property at the nodes. See Figure 2. Nevertheless, spline interpolation converges as the mesh is refined, whereas the stationary interpolation with compactly supported RBFs does not converge and exhibits error saturation. In this note we explain the reason behind this observation. We then discuss two ways in which the saturation problem can be overcome.

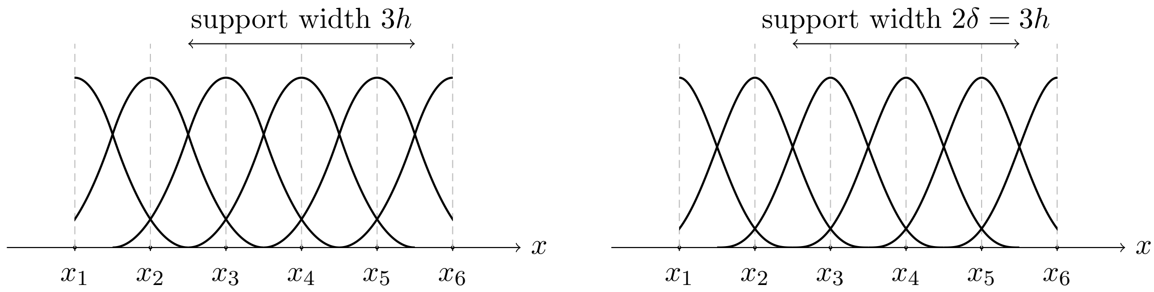

Figure 1: Quadratic B-spline basis functions (left) and scaled compactly supported Wendland RBF translates (right) on the same equidistant centers. Two families look very similar but interpolation with the basis functions on the left converges under mesh refinement, whereas interpolation with those on the right does not converge!

2 Stationary interpolation and the basic puzzle

Consider interpolation on $\Omega = [0, 1]$ with uniform points $x_j = jh$, $j = 0, \ldots, N$, where $h = 1/N$. Let ϕ be a compactly supported function satisfying $\operatorname{supp} \phi \subset [-1, 1]$ and $\phi(0) = 1$, and define the scaled function

$$\phi_\delta(x) = \phi\left(\frac{x}{\delta}\right).$$

A stationary scaling means that $\delta = ch$ with a fixed constant $c > 0$. Thus the support width shrinks at the same rate as the mesh size. At first sight, this is very similar to piecewise linear interpolation. The hat functions also have support of size $O(h)$, and the local shape of the basis does not change as the grid is refined. Hat functions are B-spline functions of degree one. The same holds true for B-spline function of higher order, for example quadratic and cubic ones. Yet the two methods may behave very differently. The main point of this article is that *locality* *and nodality do not imply convergence.* What is additionally needed is a suitable *consistency* or *reproduction property.*

To see this, it is enough to examine the constant function $f(x) \equiv 1$. Suppose first that $0 < \delta < h$. Then

$$\phi_\delta(x_i - x_j) = \delta_{ij}$$

because different supports do not reach neighboring centers. Hence the interpolation matrix is the identity and the interpolant is simply

$$I_h f(x) = \sum_j f(x_j)\phi\left(\frac{x - x_j}{\delta}\right).$$

For $f \equiv 1$, this becomes

$$I_h 1(x) = \sum_j \phi\left(\frac{x - x_j}{\delta}\right).$$

If $\delta < h/2$, gaps occur between neighboring supports. On those intervals $I_h 1(x) = 0$, and therefore $\|1 - I_h 1\|_{L^\infty(0,1)} = 1$ for every h. The interpolant is exact at every node and still does not converge. Thus the Kronecker property $\phi_j(x_i) = \delta_{ij}$ only guarantees interpolation at the nodes. It says nothing about the behavior between them.

The first consistency condition one should check is instead $I_h 1 = 1$. For a nodal basis $\{\psi_j\}$, this is equivalent to

$$\sum_j \psi_j(x) = 1$$

that is, to the partition-of-unity property.

3 Why spline interpolation converges

For simplicity we consider the splines of degree one. Let

$$B(\xi) = (1 - |\xi|)_+$$

be the standard hat function and define $B_j(x) = B((x - x_j)/h)$. The hat functions satisfy both

$$B_j(x_i) = \delta_{ij} \quad \text{and} \quad \sum_j B_j(x) = 1.$$

Even more importantly,

$$\sum_j x_j B_j(x) = x.$$

Thus all linear functions are reproduced exactly. This immediately explains the usual second-order interpolation error. For a smooth function, $f(x_j) = f(x) + f'(x)(x_j - x) + O(h^2)$. Hence

$$\sum_j f(x_j)B_j(x) = f(x)\sum_j B_j(x) + f'(x)\sum_j (x_j - x)B_j(x) + O(h^2).$$

The first sum is one and the second is zero, so

$$|f(x) - I_h^{\text{spl}} f(x)| = O(h^2).$$

Therefore linear spline convergence is not a consequence merely of compact support or nodality. The decisive fact is polynomial reproduction.

4 The stationary RBF defect

Now take $\delta = ch$ with c fixed. For the constant function,

$$I_h 1(x) = \sum_j \phi\left(\frac{x - jh}{ch}\right).$$

Introduce the normalized coordinate $\xi = x/h$. Then

$$I_h 1(h\xi) = \sum_j \phi\left(\frac{\xi - j}{c}\right).$$

Define

$$P_c(\xi) = \sum_j \phi\left(\frac{\xi - j}{c}\right).$$

The crucial observation is that P_c is independent of h. Thus, in the interior of a sufficiently large uniform grid, $I_h 1(h\xi) = P_c(\xi)$ at every refinement level. If $P_c(\xi) \not\equiv 1$, then the defect $1 - P_c(\xi)$ does not decrease as $h \to 0$. The error pattern merely becomes narrower in physical coordinates. Its amplitude stays essentially unchanged. This is the basic mechanism behind stationary saturation.

A particularly clear example is provided by the one-dimensional Wendland function

$$\phi(r) = (1 - r)_+^3 (3r + 1).$$

Take $\delta = h$. Since $\phi(1) = 0$, the interpolation matrix is again the identity. At the midpoint $x = x_j + h/2$ between two neighboring nodes, the two neighboring basis functions contribute equally, so $I_h 1(x) = 2\phi(1/2)$. Since $\phi(1/2) = (1/2)^3(3/2 + 1) = 5/16$, we obtain

$$I_h 1\left(x_j + \frac{h}{2}\right) = \frac{5}{8}.$$

Hence the midpoint error is

$$\left|1 - I_h 1\left(x_j + \frac{h}{2}\right)\right| = \frac{3}{8}.$$

This error is independent of h. Therefore

$$\|1 - I_h 1\|_{L^\infty(0,1)} \geq 3/8$$

at every refinement level. The corresponding calculation for the hat function gives $2B(1/2) = 1$, so the constant function is reproduced exactly. This single midpoint comparison already explains the difference. Both basis functions are compactly supported, both have support proportional to h, and both are nodal in this scaling. Only one has the correct reproduction property.

5 Overlapping supports and cardinal RBF bases

The previous example is especially transparent because the interpolation matrix is diagonal. In ordinary RBF interpolation, the supports may overlap, and one solves

$$s_h(x) = \sum_j c_j \phi\left(\frac{x - x_j}{\delta}\right), \qquad s_h(x_i) = f(x_i).$$

If $\delta = ch$, then $A_{ij} = \phi((i - j)/c)$. Thus, apart from boundary effects, the normalized interpolation matrix is independent of h.

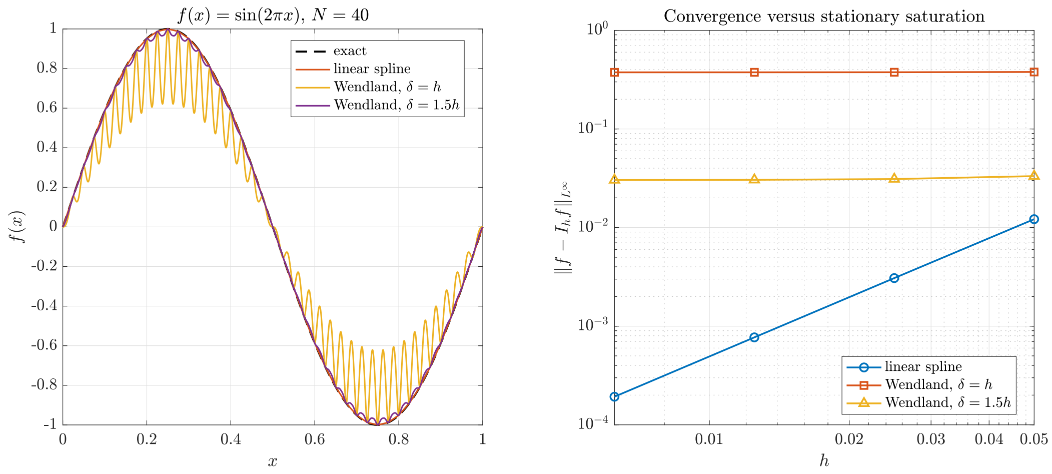

Figure 2: Comparing the exact solution with linear spline interpolation and two stationary interpolations with a compactly supported RBF of Wendland type.

The interpolant can be written in cardinal form as

$$s_h(x) = \sum_j f(x_j) L_j(x)$$

where $L_j(x_i) = \delta_{ij}$. Again, the cardinal property is not the decisive issue. The relevant question is whether

$$\sum_j L_j(x) = 1.$$

If this identity fails by a fixed amount in normalized coordinates, then refinement repeats essentially the same defect. This also explains why increasing the ratio δ/h can reduce the saturation level without necessarily eliminating it. More neighboring basis functions interact, so the stationary cardinal basis may approximate constants more accurately, but as long as the normalized defect is fixed, the error need not converge to zero.

For example, numerical experiments with the Wendland function above give approximately

δ/h	stationary L^∞ error for $f = 1$
0.4	1.000
1.0	0.375
1.5	0.043

and these values remain essentially unchanged as the mesh is refined. For the smooth function $f(x) = \sin(2\pi x)$, a the experiment results in Figure 2.

The spline error decreases by roughly a factor of four whenever h is halved (i.e. with convergence order 2), whereas the stationary RBF errors approach nonzero levels. The left plot shows that the error is large between interpolation points.

6 Polynomial reproduction and the Strang–Fix viewpoint

The phenomenon can also be interpreted through shift-invariant approximation. For a generator ϕ, consider

$$S(\phi) = \mathrm{span}\{\phi(\cdot - j) : j \in \mathbb{Z}\}.$$

Approximation order is closely related to polynomial reproduction. The classical Strang–Fix conditions describe this property in the Fourier domain [1, 4]. Roughly speaking, reproduction of constants and polynomials requires suitable zeros of $\widehat{\phi}$ and its derivatives at the nonzero reciprocal lattice frequencies $2\pi k$, $k \in \mathbb{Z} \setminus \{0\}$. For the hat function,

$$\widehat{B}(\omega) = \left(\frac{\sin(\omega/2)}{\omega/2}\right)^2$$

up to Fourier normalization, and therefore $\widehat{B}(2\pi k) = 0$ for $k \neq 0$. These zeros are closely related to the partition-of-unity and polynomial reproduction properties of the spline space.

Generic Wendland functions are designed to provide compact support, positive definiteness, smoothness, and minimal polynomial degree [5, 6]. They are not designed to satisfy lattice-dependent polynomial reproduction conditions. Thus there is no reason for their stationary translates to behave like B-spline generators.

7 How normalization changes the situation

The discussion immediately suggests a remedy. Suppose $\phi_j(x) = \phi_\delta(x - x_j)$ and define

$$\psi_j(x) = \frac{\phi_j(x)}{\sum_k \phi_k(x)},$$

assuming the denominator is positive. Then $\sum_j \psi_j(x) = 1$. The corresponding approximation

$$Q_h f(x) = \sum_j f(x_j)\psi_j(x)$$

therefore reproduces constants exactly: $Q_h 1 = 1$. If the basis is local and nonnegative, then for a Lipschitz function,

$$\begin{aligned} |f(x) - Q_h f(x)| &= \left|\sum_j \psi_j(x)\big(f(x) - f(x_j)\big)\right| \\ &\leq Ch\|f'\|_{L^\infty}. \end{aligned}$$

Thus partition of unity already restores a first-order approximation property. A closely related idea is to normalize the RBF interpolant itself:

$$R_h f(x) = \frac{I_h f(x)}{I_h 1(x)}.$$

Then $R_h 1 = 1$ exactly. Such rescaled localized RBF methods have been studied, for example, in [2]. Their behavior reinforces the main conclusion: stationary scaling is not inherently incompatible with convergence; the important issue is the reproduction structure of the stationary approximation scheme.

Higher-order convergence requires more than constant reproduction. If

$$\sum_j \psi_j(x) = 1, \quad \text{and} \quad \sum_j (x_j - x)\psi_j(x) = 0,$$

then linear functions are reproduced, and Taylor expansion leads to second-order consistency. More generally, reproduction of polynomials through degree m provides the usual mechanism for higher-order approximation.

8 How multiscale interpolation avoids stationary saturation

The saturation phenomenon described above does not mean that compactly supported RBFs cannot be used with support radii proportional to the local mesh width. An important way to recover convergence is to use the kernels in a *multiscale* or *multilevel* correction scheme rather than as a single stationary interpolant. This approach was suggested in [3] and analyzed, for example, in [7].

Let $X_1, X_2, \ldots$ be successively finer sets of centers, with fill distances $h_1 > h_2 > \cdots$, and choose at level j a scaled compactly supported kernel ϕ_{δ_j} with δ_j related to h_j. Starting with $u_0 = 0$, one does not interpolate the original function again at every level. Instead, one interpolates the current residual. More precisely, if

$$e_{j-1} = f - u_{j-1},$$

then one computes a correction

$$s_j = I_{X_j,\phi_{\delta_j}} e_{j-1}$$

and updates

$$u_j = u_{j-1} + s_j, \qquad e_j = e_{j-1} - s_j.$$

Thus the final approximation after n levels is

$$u_n = s_1 + s_2 + \cdots + s_n.$$

At first sight, this may seem puzzling. Each individual level may use a stationary-type scaling, for example $\delta_j \approx h_j$, and a one-level interpolant with this scaling may exhibit saturation. Why, then, can the multiscale method converge?

The reason is that the same stationary approximation defect is not repeated on the original function. At level j, the kernel is applied to the error left by all previous levels. The coarse levels first remove the components of the function that can be represented on large spatial scales, while the finer levels act on the remaining higher-frequency or more localized components. Thus the role of the basis functions changes from "approximate the entire function at one fixed normalized scale" to "correct the part of the error that remains at the current scale." This can be expressed schematically as

$$f \longrightarrow e_1 \longrightarrow e_2 \longrightarrow \cdots,$$

where, under suitable assumptions on the refinement ratio and the scales δ_j, one obtains a contraction

$$\|e_j\|_{\mathcal{H}_{j+1}} \le \alpha \|e_{j-1}\|_{\mathcal{H}_j}, \qquad 0 < \alpha < 1,$$

in an appropriate scale-dependent native or Sobolev norm. Iterating this estimate gives

$$\|e_n\| \le C\alpha^n \|f\|,$$

so the accumulated multilevel approximation converges even though an isolated stationary interpolation step need not converge as $h \to 0$.

The multiscale method therefore does not make the stationary one-level interpolant suddenly reproduce constants or polynomials. Rather, convergence is obtained from the *error-correction mechanism across scales*. The approximation space after several levels is the sum

$$V_1 + V_2 + \cdots + V_n,$$

which is considerably richer than any individual stationary space V_j. In this sense, the missing approximation power of a single scale is supplied collectively by the hierarchy of scales.

This viewpoint is closely related to multiresolution ideas. Coarse levels capture large-scale features, finer levels correct what the coarse approximation misses, and the scale-dependent residual becomes the quantity that contracts from one level to the next. This is the basic reason why multiscale interpolation with compactly supported kernels can overcome the saturation encountered in ordinary one-level stationary interpolation.

9 Concluding remarks

The difference between stationary compactly supported RBF interpolation and linear splines is not explained by locality or nodality alone. Compact support provides locality, while the Kronecker property guarantees interpolation at the data sites, but neither ensures convergence. Linear hat functions additionally reproduce $\mathbb{P}_1$. Compactly supported RBFs do not possess such reproduction properties.

For a stationary RBF scaling $\delta = ch$, the local approximation remains unchanged in normalized coordinates as $h \to 0$. Consequently, a fixed reproduction defect is repeatedly compressed into smaller intervals without decreasing in amplitude. This leads to saturation. This difficulty can be overcome in different ways. Normalization, polynomial augmentation, and related techniques introduce the missing reproduction properties. Multiscale interpolation follows a different strategy: at each level the scaled kernel is applied to the residual left by the previous levels. The resulting approximation belongs to the increasingly rich space $V_1 + \cdots + V_n$, and under suitable conditions the residual contracts from one level to the next.

Thus, stationary scaling itself does not prevent convergence. What matters is whether sufficient approximation power is provided either in each stationary space through consistency and polynomial reproduction, or across a hierarchy of spaces through multiscale residual correction.

References

[1] M. D. Buhmann, N. Dyn, and A. Ron. Compactly supported radial functions and the Strang–Fix condition. *Applied Mathematics and Computation*, 84(2–3):115–124, 1997.

[2] S. De Marchi and H. Wendland. On the convergence of the rescaled localized radial basis function method. *Applied Mathematics Letters*, 99:105996, 2020.

[3] R. Schaback. Creating surfaces from scattered data using radial basis functions. In Morten Dæhlen, Tom Lyche, and Larry L. Schumaker, editors, *Mathematical Methods for Curves and Surfaces*, pages 477–496. Vanderbilt University Press, Nashville, TN, 1995.

[4] G. Strang and G. J. Fix. A Fourier analysis of the finite element variational method. In G. Cremonese, editor, *Constructive Aspects of Functional Analysis*, volume II, pages 793–840. Edizioni Cremonese, Rome, 1973.

[5] H. Wendland. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in Computational Mathematics*, 4(1):389–396, 1995.

[6] H. Wendland. *Scattered Data Approximation.* Cambridge University Press, Cambridge, 2005.

[7] H. Wendland. Multiscale analysis in Sobolev spaces on bounded domains. *Numerische Mathematik*, 116(3):493–517, 2010.