Snapshots of modern mathematics
from Oberwolfach

№ 2/2019

On radial basis functions

Martin Buhmann • Janin Jäger

Many sciences and other areas of research and applications from engineering to economics require the approximation of functions that depend on many variables. This can be for a variety of reasons. Sometimes we have a discrete set of data points and we want to find an approximating function that completes this data; another possibility is that precise functions are either not known or it would take too long to compute them explicitly. In this snapshot we want to introduce a particular method of approximation which uses functions called *radial basis functions*. This method is particularly useful when approximating functions that depend on very many variables. We describe the basic approach to approximation with radial basis functions, including their computation, give several examples of such functions and show some applications.

1 Why do we need approximation?

In many scientific and engineering applications, we find examples of functions that depend on very many unknown quantities, or *variables*. It is easy to imagine functions in physical sciences that depend on the measurement of a point in space $p = (x, y, z)$ and the time t at which the measurement was taken, for instance, the function that monitors global temperature over time. Then the temperature is a function $f(x, y, z, t)$ of four unknowns. For an example from economics, the total price function of a collection of investments depends on the price of all the individual shares making up the portfolio. Sometimes, theoretically and in applications, we find functions that depend on literally hundreds and thousands of variables. These functions are often complicated and "expensive" to evaluate, which means that the computation of the value of the function at a given point takes a long time and a lot of computing power. Especially when a very accurate evaluation is needed, this is a big problem. Another difficulty arises when the values of the function are unknown, sometimes all that is known is that the function exists. This situation is typical when solving "differential equations", where a differential equation is an equation that involves a function and its derivatives. In applications, the function can represent some physical quantity and the derivative is the rate of change of this quantity. These types of relations are common, so differential equations, and hence the need to approximate solutions, arise often in mathematics, engineering, physics, biology and so on.

Computation can also be difficult when there is a lot of data to handle. Returning to the global temperature example, if we want to make a weather forecast from the data collected from the weather stations shown in Figure 1, knowing the temperature values only at the stations is not sufficient. We would need to find an approximation for the temperature values also in coordinates between the weather stations to produce a map. Notice that the temperature will depend on at least two variables if we look at the landscape as a two dimensional map, and, furthermore, the location of the weather stations is unchangable - we will see shortly that this could be a problem.

To summarise, we need to find ways to approximate functions, which are either too complicated to calculate explicitly or are unknown at some or all

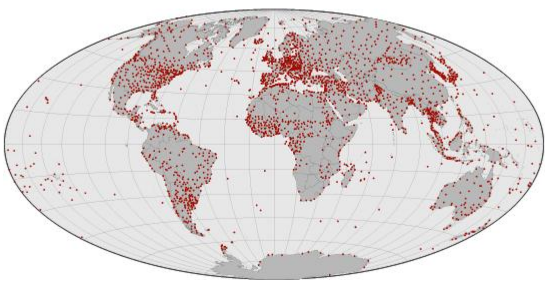

Figure 1: Weather stations on the surface of the earth.

coordinates, by simpler functions, subject to conditions such as accuracy and a limit on the number of computer operations to evaluate them within a certain time-frame.

1.1 What is approximation?

When approximating we try to get from an unknown or complicated function f to a known and easy to evaluate function s. Moreover, this s should be "close" to the original function. In the weather example, the real temperature function f is only known at the weather stations, so we try to find an approximating function s to the temperature which can be easily evaluated at any point of the map.

There are many ways to find such approximating functions, but to choose an appropriate one we must first decide how the original function f and its approximated version s should be related. If we want them to match exactly at certain points, very frequently we would use *interpolation*. Roughly speaking, and we will explain in more detail below, interpolation in a mathematical sense means any method of constructing new data points within a range of finitely-many given data points. So, we know the values of a function f at a given set of points, and the approximating function s obtained by interpolation from these points has to match these values (like the temperature values at the weather stations). This is not the only option. In some situations the values of the derivative of f at certain points may be known, and we may ask additionally that the approximating function s has derivatives that match those of f.

In other situations interpolation makes less sense. For instance, in physical measurements there are often little errors in the measured data, sometimes they are produced by the equipment and sometimes by the environment. These errors are called *noise* and we would not want our approximation to reproduce the noise. In other words, given that the original function is not measured absolutely accurately, it is not important to have these data points match exactly with the values taken by the approximating function.

The most serious problem for interpolation is that sometimes an approximation simply cannot be obtained in this way. We will show how approximation in the absence of good conditions for interpolation can be done in the last section of the snapshot, but first, in the following section, we will explain interpolation in more detail.

1.2 A well known interpolation: using polynomial functions

In one dimension, interpolating functions always exist. In fact, *polynomials* can perform this type of approximation if their *degree* is sufficiently high. We recall

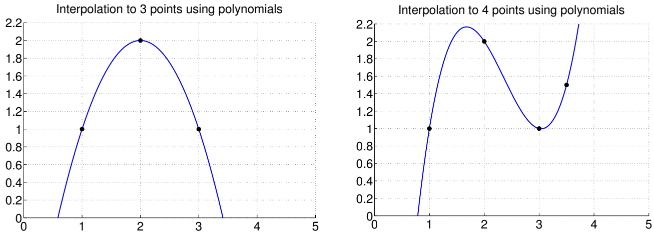

Figure 2: Polynomial-Interpolation to 3 and 4 points

that a polynomial is a function with the form

$$s(x) = a_n x^n + a_{n-1} x^{n-1} + \cdots + a_1 x + a_0,$$

where $a_0, a_1, \ldots, a_n$ are real numbers called the *coefficients* of s, and, supposing that a_n is not equal to zero, the degree of the polynomial is n.

We start here by giving a simple example of how polynomial interpolation works. If we want to find an approximate version of a function $f : \mathbb{R} \to \mathbb{R}$ and we know that $f(1) = 1$, $f(2) = 2$ and $f(3) = 1$, then we can choose a polynomial $p(x) = ax^2 + bx + c$ of degree 2 that interpolates f at the known points, where the coefficients a, b and c are determined by solving the system of three linear equations $p(1) = f(1) = 1$, $p(2) = f(2) = 2$ and $p(3) = f(3) = 1$. The solution turns out to be $a = -1$, $b = 4$ and $c = -2$. If we add the point $f(3.5) = 1.5$, we can choose a polynomial $p(x) = ax^3 + bx^2 + cx + d$ of degree three; in this case we must solve a system of four linear equations to determine the approximating function. The graphs of both cases are shown in Figure 2. One can see that to include more points, we need higher order polynomials. Also, the approximate function will always tend to ∞ or $-\infty$ for big and small values of x, which is something we usually do not want to happen, given that it is not very realistic when we have a set of physical measurements like in the weather example. This is one simple reason for polynomial interpolation not always being useful, but there is also a deeper reason that we address below[1].

1.3 Problems with more dimensions

In the applications already mentioned, we rarely use data that depends on only one variable. Rather, the points at which we know the value of the function f come from a space of at least two dimensions. When the dimension of this

[1] To read more about interpolation, see the previous snapshot https://imaginary.org/snapshot/a-few-shades-of-interpolation.

background space is two or more, there always exist distributions of data sites for which we cannot find a polynomial that interpolates the data. To make this more precise, we need to introduce another definition. We want to find an approximating function from a set of functions (in the example above we used the set of all polynomials, but we can also choose another set). These functions can be built up using *basis functions*, which are a collection of "simpler" functions chosen such that every function in our set can be written as a *linear combination* of these basis functions, where a linear combination is a sum of basis functions multiplied by real numbers. For the set of polynomial functions of one variable that we introduced above, the basis functions could be $\{1, x, x^2, x^3, \ldots\}$; this collection is the "usual" choice, but there exist others. Then, a theorem due to Mairhuber [6] says that if we fix a set of basis functions and attempt to interpolate arbitrary data points in a space of two or more dimensions, there will always exist a spatial distribution of these data points for which interpolation using a function from the set with the given basis is impossible. Another way to say this is that the basis functions must be chosen after we know the data points. This significantly increases the amount of computation needed.

1.4 Other approximation methods

There are many other ways to make approximations. Instead of matching points and filling in the gaps like we have done with interpolation functions, we could try instead to find approximations that minimise the overall "distance" between the original function and its approximation. An example is the *least-squares approximation*, where the aim is to find an approximating function that has the property that the sum of the squares of the distances between known data points and those predicted by the approximation is as small as possible[2]. For example, if we start with the three points $(0, 6), (1, 0)$ and $(2, 0)$ in the plane, there is no straight line that passes through all three of these points (to see why, try drawing a diagram). However, we can ask for the straight line that passes as close as possible in the sense of least-squares. This line turns out to be $y = 5 - 3x$.

There is also no need to only use polynomials. Approximation can be done using *piecewise* polynomials, also known as *splines*, where there are many pieces of different polynomials joined at certain points with conditions on how they can be fitted together. There are also other functions such as exponentials, or even piecewise exponentials, known as *exponential splines*, or trigonometric functions which are well suited to be used for interpolation or other ways to approximate one-dimensional data. But we are still left with the problem that

[2] To read more about the least-squares method, see for instance, https://en.wikipedia.org/wiki/Least_squares.

approximating functions do not necessarily exist for all possible distributions in space of data points, just as in the polynomial interpolation case outlined above, as soon as the dimension is more than one. Polynomials, splines, and so on, are all defined in multiple dimensions, but the existence of sufficiently easy to compute approximating functions is not at all trivial. Therefore the question naturally arises: What should we do in two, three or indeed 10000 dimensions? This question is also relevant even if we do not use interpolation techniques, because the basics of the approximation process will still depend on basis functions and understanding the functions that can be built using them.

In the following section, we investigate a method of approximation that avoids this dimension problem: that of *radial basis functions*.

2 Approximation using radial basis functions

Generally speaking, radial basis functions are means to approximate continuous functions that depend on two or more variables (we will refer to them as *multivariate* functions) by linear combinations of terms based on a single function of one variable. We will denote this function by $\phi : \mathbb{R}^+ \to \mathbb{R}$ and call it a *radial basis function* [4]. Note the difference: now we have a single basis function with which to build our approximations, before we had a whole set of basis functions.

To get from ϕ to a multivariate function, we define a new function Φ as follows: For each $\mathbf{x} = (x_1, \ldots, x_n)$ in $\mathbb{R}^n$, we set

$$\Phi(\mathbf{x}) = \phi(\|\mathbf{x}\|_2).$$

Here $\|\mathbf{x}\|_2$ is the Euclidean distance of the point $\mathbf{x}$ from the origin, which is computed as $\|\mathbf{x}\|_2 = \sqrt{x_1^2 + x_2^2 + \cdots + x_n^2}$. In the Euclidean plane, that is, in two dimensions, if we fix a circle with centre at the origin and radius r, and if the point (x_1, x_2) lies on this circle, then $\|(x_1, x_2)\|_2 = r$. In other words, for every point on a circle centred at the origin, the function Φ has the same value. The same is true for spheres centred at the origin in three-dimensional space, and their higher-dimensional equivalents. So we see that function $\Phi(\mathbf{x})$ only depends on the distance of each point $\mathbf{x}$ from the origin. In technical terms, we say it is *radially symmetric*. This is the reason for the name “radial basis function”.

Let us consider a specific example, the *inverse multiquadric*. The inverse multiquadric has the form

$$\phi(x) = \frac{1}{\sqrt{1 + x^2}}$$

and is shown on the left-hand side of Figure 3. On the right-hand side of the same figure, we have replaced x by the two-dimensional Euclidean distance, to obtain the function $\Phi((x_1, x_2)) = \frac{1}{\sqrt{1+(x_1^2+x_2^2)}}$. This is the two-dimensional inverse multiquadric.

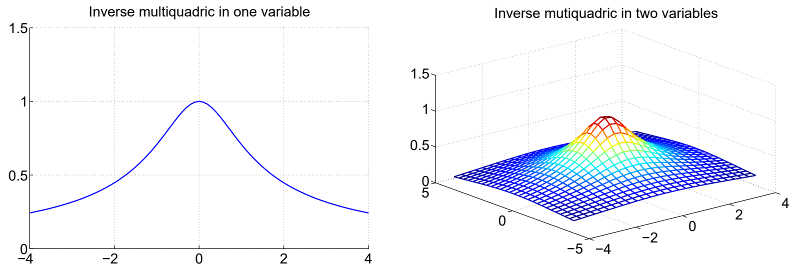

Figure 3: On the left, the inverse multiquadric basis function $\phi(|x|)$, and, on the right, the two-dimensional function $\Phi(\mathbf{x})$.

It is easy to compute such functions in higher dimensions, even though we are not able to visualize their graphs. Later we will discuss more examples of radial basis functions, but first we will find out how they are used.

2.1 Interpolation using radial basis functions

When we want to perform an interpolation to a function $f : \mathbb{R}^n \to \mathbb{R}$ which is known at a finite number of points $x_j \in \mathbb{R}^n$, which from here on we will also call *centres*, as already discussed above, we need a set of basis functions. In the polynomial case, these were 1, x, x^2 and so on, which we combined to produce the approximating function s. Now we produce a set of basis functions using our single radial basis function ϕ by shifting it by the given data points. In this way, the function $\phi(\|x - x_j\|_2)$ has the same shape as the radial function Φ but shifted to the point x_j instead of the origin. So for every centre, we get one new basis function and we will use these to form an approximating function s. This s will be a linear combination of the basis functions, in other words it has the form

$$s(x) = \sum_j \lambda_j \phi(\|x - x_j\|_2),$$

where each λ_j is a real number. We want to choose s to interpolate f at the given data points, meaning the conditions $s(x_k) = f(x_k)$ should hold for all k. The real coefficients λ_j can then be computed by solving a system of linear equations arranged in a matrix, which we call the *interpolation matrix*.

The method is best explained while computing a little example. We consider the same interpolation problem we solved before using polynomials. Recall that the function f is known to have the values $f(1) = 1$, $f(2) = 2$ and $f(3) = 1$. Let us now interpolate using the inverse multiquadric function described above, that is, as basis functions we will use shifted copies of the radial basis function $\phi(x) = \frac{1}{\sqrt{1+x^2}}$ with centres in $x_1 = 1$, $x_2 = 2$, $x_3 = 3$. The basis functions are shown in the first picture of Figure 4. The approximating function s is a linear combination of these shifted copies, so has the form

$$s(x) = a\phi(|x-1|) + b\phi(|x-2|) + c\phi(|x-3|).$$

Just as in the polynomial case, we must find a, b and c in such a way that the equations $s(1) = f(1) = 1$, $s(2) = f(2) = 2$ and $s(3) = f(3) = 1$ hold. Therefore,

$$s(1) = a\phi(|1-1|) + b\phi(|1-2|) + c\phi(|1-3|) = 1,$$

and then, computing the values $\phi(|1-1|) = \phi(0) = \frac{1}{\sqrt{1+0}} = 1$, $\phi(|1-2|) = 1/\sqrt{2}$ and $\phi(|1-3|) = 1/\sqrt{5}$, we simplify this equation as follows:

$$s(1) = a \cdot 1 + b \cdot \frac{1}{\sqrt{2}} + c \cdot \frac{1}{\sqrt{5}} = 1.$$

Performing the same simplification with the other two conditions (that is, $s(2) = 2$ and $s(3) = 1$) we obtain the system of three linear equations:

$$A := \phi(|x_i - x_j|)_{i,j} = \begin{pmatrix} 1 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{5}} \\ \frac{1}{\sqrt{2}} & 1 & \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{5}} & \frac{1}{\sqrt{2}} & 1 \end{pmatrix}.$$

This is the interpolation matrix. It is symmetric, not only in our example but always, by its definition. To compute the values of a, b and c, we either solve the system of linear equations, or we calculate the inverse of the matrix and multiply it by the vector on the left side of the equation

$$\begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix} = A \cdot \begin{pmatrix} a \\ b \\ c \end{pmatrix}.$$

This shows that invertibility of the interpolation matrix and solvability of the interpolation problem are equivalent in this context.

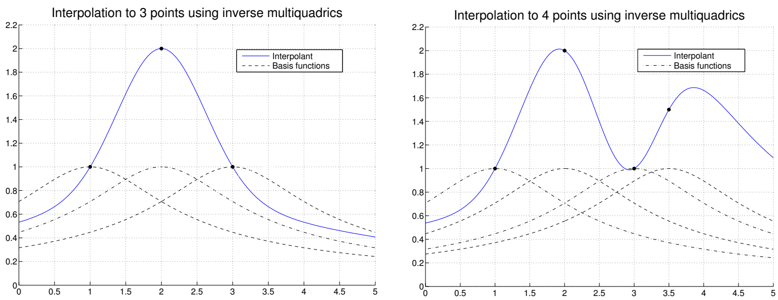

Figure 4: Interpolation using Inverse Multiquadrics to 3 and 4 points

The solution s is shown in Figure 4, where we also see the solution we get when we include the point $f(3.5) = 1.5$. The figure highlights the fact that the additional basis function $\phi(|x - 3.5|)$ is dependent on the data site $x_4 = 3.5$. This *data-dependence* is the big advantage that radial basis functions have over the polynomial interpolation, because it makes interpolation possible in higher dimensions.

A particularly important feature of this approach is that it provides *global* approximations on $\mathbb{R}^n$ that are highly accurate under suitable conditions on the functions f which are to be approximated.

2.2 Examples and application

One of the greatest further advantages of this method lies in its adaptivity to the function to be interpolated. An important property to think about is the "smoothness" of the function we want to approximate. For example, we might think of temperature as a differentiable function, but if we model the price of a share on the stock market, differentiability is not wanted. The reason is that stock prices are influenced by a variety of factors in the short term (sales of the firm, launch of new products, political decisions, for instance elections). These factors lead to sudden changes to the price at a certain point in time, and these sudden changes mean that the price function is not smooth, that is, it is not differentiable. The approximation inherits the smoothness of the basis function, so we can choose the basis function we use according to what properties we think (or know) that the approximation enjoys. Let us look at some other common radial basis functions.

A typical example of a radial basis function is the multiquadrics function with a real parameter c [7]

$$\Phi(x) = \phi(\|x\|_2) = \sqrt{\|x\|_2^2 + c^2}, \qquad x \in \mathbb{R}^n.$$

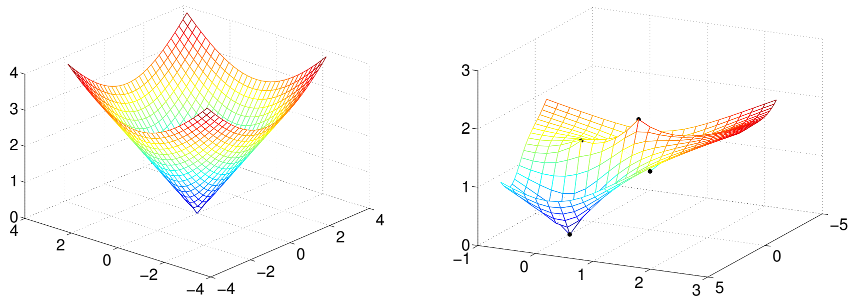

Figure 5: The linear basis function $\Phi(x) = \|x\|_2$ and an interpolation performed using it in two dimensions.

Further examples are reciprocal multiquadrics $\phi(x) = 1/\sqrt{x^2 + c^2}$ and the *Gauss-kernels* $\phi(x) = \exp(-c^2 x^2)$, which are often used in probability theory. All of these are differentiable functions and therefore they result in differentiable approximating functions.

Alternatively, we could use the linear basis function

$$\Phi(x) = \|x\|_2,$$

which is shown in Figure 5, which is continuous but not differentiable (notice the "corner" at the origin), and so produces only continuous approximating functions. For a two-dimensional interpolation obtained using this linear radial basis function, see the right-hand side of Figure 5.

All the above-mentioned radial basis functions lead to what are called "positive definite" interpolation matrices and therefore they are called positive definite functions [1, 2, 7]. Performing interpolation in an n-dimensional space using a positive definite function, there are no restrictions on the data, neither in amount nor in distribution, except that they need to be distinct points, so that it will always be possible to find an interpolating function of the form described. This should be contrasted to polynomial interpolation where data points being distinct is a necessary but not sufficient condition for the uniqueness and existence of interpolating functions. A further advantage is that the computational effort is low [1, 8].

2.3 A meshfree method

For applications it is desirable, especially in high dimensions, that little extra work is necessary before the data points can be used. Fortunately this is the case for radial basis function algorithms, whereas for instance "finite element" or multivariate spline methods – or indeed any piecewise defined approximation scheme – normally need triangulations or analogous constructions in higher dimensions. In fact, the advance structuring of the data that some other approximation schemes depend on can be prohibitively expensive to compute. Therefore our approximations here are considered to be *meshfree* approximations [3].

On the other hand, if we start with a large data set, advanced numerical methods called *fast algorithms* for computing the radial basis function approximations can be required, due to the radial basis functions' frequent unboundedness. If the data set is small enough, we only need standard software.

3 Beyond interpolation: What is Quasi-Interpolation?

When we are working with perhaps thousands of data sites, solving the linear equations to compute an interpolating function can take a lot of time and computing capacity. To solve this problem, a new method called *quasi-interpolation* was introduced to perform approximation without having to solve those equations.

We take one basis function Ψ that is close to one at the origin and decays rapidly to zero away from the origin. This function is then shifted in the same way as described above for interpolation (see Figure 6, first picture). The shift with the centre at x_j is then multiplied by the function f's value at x_j (as shown in the second picture of Figure 6). Summing up the so-obtained basis functions yields the approximation

$$s(x) = \sum_i f(x_i)\Psi(x - x_i),$$

as shown in the third picture in Figure 6.

We observe that many of the radial basis functions, such as the multiquadrics, do not decay away from the origin in the required way. Perhaps surprisingly, this is not a problem at all, as linear combinations of them which have all the necessary good properties can be used instead, and they are often fairly simple to compute. For an example, consider the multiquadric $\phi(x) = \sqrt{c^2 + \|x\|^2}$, shown on the left of Figure 7, which as it is cannot be used for quasi-interpolation. However, if we compute

$$\Psi(x) = -\frac{1}{2}\phi(x-1) + \phi(x) - \frac{1}{2}\phi(x+1),$$

which is shown on the right side of Figure 7, we see that it can be used for quasi-interpolation. In fact, this is the function that was used to compute the

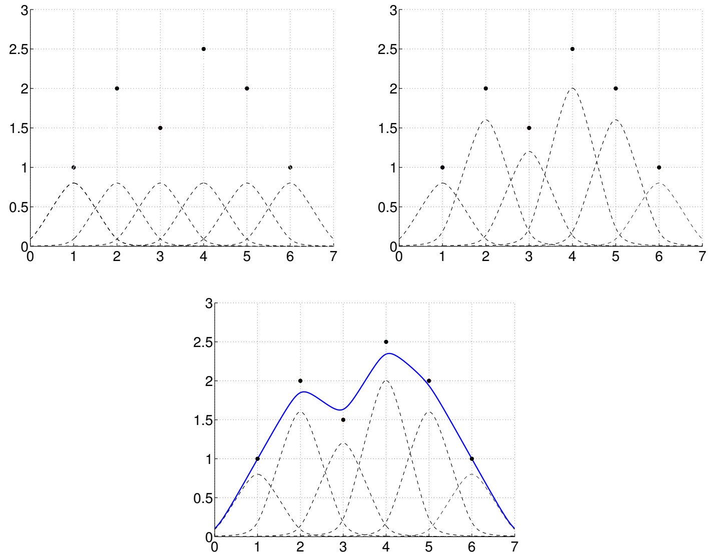

Figure 6: Quasi-Interpolation performed using the multiquadric from Figure 7

approximation in Figure 6. This is a very simple example and it needs to be said that finding quasi-interpolants in higher dimensions and with scattered data points is not always so easily done.

Recent research on quasi-interpolation has shown that, for example, quasi-interpolation using multiquadrics, inverse multiquadrics and other examples in n dimensions to "smooth enough" functions achieve very small errors, so long as the centres are spaced in a regular enough way [1]. Quasi-Interpolation is an area where there is a lot of research to be done. Just recently, quasi-interpolation was adapted for scattered data in high dimensional spaces. There is also current research on how to adapt quasi-interpolation to the surface of a sphere, which is useful when we have data from the surface of the earth to approximate.

4 Recent developments

In recent work, not yet published, the authors generalise the concepts of radial basis functions so as to be able to provide unique interpolants to large classes of what are called "multiply monotone" functions. This approach provides sufficient conditions on the univariate functions ϕ so that the interpolation matrices are invertible. This then allows interpolation to any set of scattered

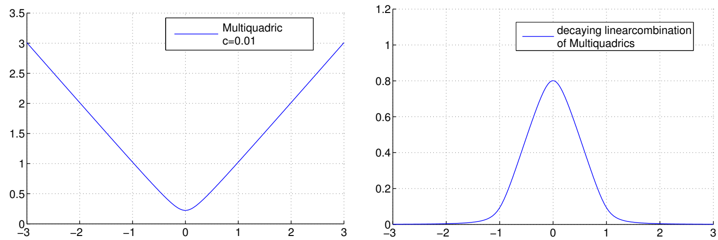

Figure 7: Multiquadric ϕ and linear combination of multiquadrics Ψ

data in high dimensional spaces to be *regular*, which means that approximating functions exist uniquely, subject to some reasonable conditions.

These theorems require the radial basis function and a fixed, finite number of its derivatives to be monotonic (that is, increasing or decreasing functions), and also that at each stage of differentiation there is a sign change. Therefore, in particular, less smooth functions (such as truncated power functions for instance, or piecewise exponentials) than the infinitely differentiable ones like multiquadrics can now be included. The above-mentioned regularity of the scattered data interpolation problem is guaranteed up to a certain upper bound on the dimension of the underlying data (but this bound can be adjusted depending on how many derivatives of the radial function are assumed to be monotonic, and if the function is infinitely differentiable and all its derivatives are monotone, there is no bound to the dimension anymore).

The conditions are not necessary, but only sufficient, in the sense that there are other ways to have regular interpolation (for example, using oscillating radial basis functions), but they are very useful for three reasons. Firstly, they are particularly straightforward to verify since they are conditions that depend on only one variable, although the resulting regularity statements apply to several dimensions. Secondly, almost all radial basis functions currently in use are included. Finally, they are closely related both to the important structural properties of the interpolation matrix and to the positivity of the Fourier transform of the function.

In another paper [5], the authors applied radial basis function interpolation to an interpolation problem on the sphere. The interpolation is used in that article to reconstruct data that were collected using an EEG (electroencephalogram). The EEG records the electrical activity of the brain using electrodes which are connected to the scalp. The radial basis function method is especially useful in this context because it can be used to derive an estimate of the activity at any point on the surface of the head from the values measured at the electrodes.

In addition, interpolation (or, in fact, extrapolation) algorithms can be used for the reconstruction of missing data which were lost due to technical problems, which might be caused by a broken electrode or some physiological artefact. In [5], we investigate whether the application of radial basis functions has advantages as compared to the more commonly used methods, and whether taking into account the special geometry of the sphere (this leads to a technique called spherical basis function method) gives an advantage over usual radial basis function interpolation, and we were able to conclude that this method is indeed a good alternative to existing methods.

Image credits

Fig. 1 NASA map by Robert Simmon, based on data from the National Climatic Data Center, Public Domain, http://earthobservatory.nasa.gov/Features/GISSTemperature/Images/global_weather_stations_map.gif, accessed on 3 August 2016.

References

[1] M. Buhmann, *Radial basis functions*, Cambridge University Press, 2003.

[2] N. Dyn, *Interpolation and approximation by radial and related functions*, Approximation Theory VI **1** (1989), 211–234.

[3] G. R. Fasshauer, *Meshfree approximation methods with matlab*, Interdisciplinary mathematical sciences, vol. 6, World Scientific Publishing, 2007.

[4] R. Franke, *Scattered data interpolation: Test of some methods*, Mathematics of Computation **38** (1982), no. 157.

[5] J. Jäger, A. Klein, M. Buhmann, and W. Skrandies, *Reconstruction of electroencephalographic data using radial basis functions*, Clinical Neurophysiology **127** (2016), no. 4, 1978–1983.

[6] J. C. Mairhuber, *On Haar's theorem concerning Chebychev approximation problems having unique solutions*, Proceedings of the American Mathematical Society **7** (1956), no. 4, 609–615.

[7] C. A. Micchelli, *Interpolation of scattered data: Distance matrices and conditionally positive definite functions*, Constructive Approximation **2** (1986), 11–22.

[8] M. J. D. Powell, *The theory of radial basis function approximation in 1990*, Advances in Numerical Analysis 2: Wavelets, Subdivision, and Radial Functions **2** (1992), 371–382.

Martin Buhmann *is a professor of numerical analysis at the Justus-Liebig University of Giessen.*

Janin Jäger *is a PhD student in numerical analysis at the Justus-Liebig University of Giessen.*

Mathematical subjects
Numerics and Scientific Computing

Connections to other fields
Computer Science, Engineering and Technology, Physics

DOI
10.14760/SNAP-2019-002-EN

Snapshots of modern mathematics from Oberwolfach provide exciting insights into current mathematical research. They are written by participants in the scientific program of the Mathematisches Forschungsinstitut Oberwolfach (MFO). The snapshot project is designed to promote the understanding and appreciation of modern mathematics and mathematical research in the interested public worldwide. All snapshots are published in cooperation with the IMAGINARY platform and can be found on www.imaginary.org/snapshots and on www.mfo.de/snapshots.

ISSN 2626-1995

Mathematisches Forschungsinstitut
Oberwolfach gGmbH
Schwarzwaldstr. 9 – 11
77709 Oberwolfach
Germany

Director
Gerhard Huisken

Mathematisches Forschungsinstitut Oberwolfach

Member of Leibniz Association

IMAGINARY open mathematics