Approximation error maps

A. Gomide and J. Stolfi

Institute of Computing, University of Campinas, Brazil
{anamaria,stolfi}@ic.unicamp.br

Abstract

In order to analyze the accuracy of a fixed, finite-dimensional approximation space which is not uniform over its domain Ω, we define *approximation error map*, a description of how the error is distributed over Ω—not for a single test function but for a general class of such functions. We show how to compute such a map from the best approximations to an orthonormal basis of the target function space.

1. Introduction

The expected accuracy of a finite-dimensional approximation space (e.g. a polynomial spline space, or a finite wavelet decomposition) will often vary over its domain Ω. Indeed, adaptive-resolution schemes are based on the premise that refining the element grid in a particular region of Ω will improve the approximation accuracy in that region.

Knowledge of how the expected approximation error varies over the domain Ω is obviously relevant to the evaluation of an approximation space, and to the tuning of knot locations, grid geometry, refinement thresholds and other parameters. Towards that goal, we introduce the concept of *approximation error map*, a description of how the error is distributed over Ω—not for a single test function, but for all functions in some specified space $\mathcal{F}$. We then show how to compute such a map from the best approximations to an orthonormal basis of $\mathcal{F}$.

1.1. Notation and definitions. Let $\mathcal{F}$ and $\mathcal{A}$ be two fixed, finite-dimensional vector spaces, not necessarily disjoint, of functions defined on some domain Ω with values in $\mathbf{R}$. Let $\|\cdot\|$ be a vector semi-norm for the space $\mathcal{A}+\mathcal{F}$. For any function $f \in \mathcal{F}$, we define its *best approximation* as the function $f^{\mathcal{A}} \in \mathcal{A}$ that minimizes the error $\left\|f - f^{\mathcal{A}}\right\|$.

We refer to $\mathcal{A}$ and $\mathcal{F}$ as the *approximation* and *gauge spaces*, respectively. We assume that the $\|\cdot\|$-balls *in the subspace* $\mathcal{A}$ are strictly convex, ensuring that the best approximation always exists and is unique. Since $(\alpha f)^{\mathcal{A}} = \alpha(f^{\mathcal{A}})$ and $\|\alpha f\| = |\alpha|\,\|f\|$ for any real constant α, we can confine the analysis of approximation errors to the *unit $\mathcal{F}$-sphere* $\mathcal{F}_1 = \{\, f \in \mathcal{F} \,:\, \|f\| = 1 \,\}$.

1.2. Global error measures. Usually, the effectiveness of the approximation space $\mathcal{A}$ is measured by a single number $\|f - f^{\mathcal{A}}\|$—either for the worst-case function $f \in \mathcal{F}_1$, or by the root-mean-power average over all functions $f \in \mathcal{F}_1$

$$\sigma^*_{p,\mathcal{A},\mathcal{F}} = \left[\int_{\mathcal{F}_1} \|f - f^{\mathcal{A}}\|^p \, df\right]^{1/p} \Bigg/ \left[\int_{\mathcal{F}_1} 1 \, df\right]^{1/p}. \tag{1}$$

Note that integrals are taken over the function space $\mathcal{F}_1$, not over the domain Ω. The worst-case error is the limit

$$\mu^*_{\mathcal{A},\mathcal{F}} \;=\; \lim_{p \to +\infty} \sigma^*_{p,\mathcal{A},\mathcal{F}} \;=\; \sup\left\{ \|f - f^{\mathcal{A}}\| \;:\; f \in \mathcal{F}_1 \right\}. \tag{2}$$

1.3. Uniform approximation spaces. A global error measure such as $\mu^*_{\mathcal{A},\mathcal{F}}$ or $\sigma^*_{p,\mathcal{A},\mathcal{F}}$ is generally sufficient when all points of Ω are equivalent with respect to the quality of approximation. More formally, we say that a normed function space $\mathcal{X}$ is *uniform* over Ω if there is some family Φ of maps from Ω to Ω that preserves $\mathcal{X}$ and its norm $\|\cdot\|$, and which can take any point of Ω to any other point. A natural example is $\mathcal{Y}^d_n$, the set of all harmonic functions on the sphere $\mathbf{S}^d$ of a given maximum order n, with any L_p norm; this space is preserved by the family of rigid rotations of $\mathbf{S}^d$. Obviously, if both $\mathcal{A}$ and $\mathcal{F}$ are uniform under the same family Φ, then $\mathcal{A}$ approximates $\mathcal{F}$ equally well at all points of Ω. (Of course, for any *specific* function $f \in \mathcal{F}$, the error $f - f^{\mathcal{A}}$ will usually vary over Ω.)

There are however many important approximation spaces $\mathcal{A}$ which are not uniform. A familiar example is the space of polynomials or trigonometric series defined on a bounded region $\Omega \subseteq \mathbf{R}^n$. Another example is the space of the piecewise polynomial splines of fixed order and continuity defined over a fixed grid G. Wavelet spaces truncated to a fixed order provide yet another example. For such spaces, the expected approximation error usually varies over Ω, even when the functions to be approximated are drawn from a uniform space.

2. Approximation error map

We define the *root mean power approximation error map of $\mathcal{F}$ by $\mathcal{A}$* as the function $\sigma_{p,\mathcal{A},\mathcal{F}}$ of Ω to $\mathbf{R}$ defined by

$$\sigma_{p,\mathcal{A},\mathcal{F}}(x) = \left[\int_{\mathcal{F}_1} \left|f(x) - f^{\mathcal{A}}(x)\right|^p \, df\right]^{1/p} \Bigg/ \left[\int_{\mathcal{F}_1} 1 \, df\right]^{1/p}. \tag{3}$$

As before, integrals are taken over the function space $\mathcal{F}_1$, not over the domain Ω. Note that $\sigma_{p,\mathcal{A},\mathcal{F}}(x)$ is not the error for a *specific* function f, but rather the *average* error *at the point x* for a *generic* function f in $\mathcal{F}_1$. As a limiting case, we define also the *worst-case approximation error map of $\mathcal{F}$ by $\mathcal{A}$* as the function

$$\mu_{\mathcal{A},\mathcal{F}}(x) \;=\; \lim_{p \to +\infty} \sigma_{p,\mathcal{A},\mathcal{F}}(x) \;=\; \sup\left\{ \left|f(x) - f^{\mathcal{A}}(x)\right| \;:\; f \in \mathcal{F}_1 \right\}. \tag{4}$$

Again note that the supremum is taken over $\mathcal{F}_1$, not over Ω, and that $\mu_{\mathcal{A},\mathcal{F}}(x)$ is not the error at x for a *single* function f, but rather the error for the function f in $\mathcal{F}_1$ that is worst *for that particular* x. A plot of $\sigma_{p,\mathcal{A},\mathcal{F}}(x)$ or $\mu_{\mathcal{A},\mathcal{F}}(x)$ over Ω should show at a glance how well $\mathcal{A}$ approximates $\mathcal{F}$ in different parts of the domain, for all functions of $\mathcal{F}$ at once.

3. Computing the approximation error map

Formulas (3)–(4) become more tractable when the function metric $\|\cdot\|$ is the L_2 norm $\|f\| = [\int_\Omega |f(x)|^2\, dx]^{1/2}$ defined on the space $\mathcal{A}+\mathcal{F}$—in other words, when $\|f\|^2 = \langle f, f\rangle$ where $\langle f, g\rangle = \int_\Omega f(x)g(x)\, dx$. We make this assumption in the remainder of this section. In that case, $f^{\mathcal{A}}$ is a linear function of f, namely the orthogonal projection of f onto the subspace $\mathcal{A}$; and $\mu_{\mathcal{A},\mathcal{F}}$ is simply $|\sin\theta|$, where θ is the angle between the two subspaces.

3.1. Explicit formula for σ. Let us suppose that $\mathcal{A}$ and $\mathcal{F}$ are disjoint, and let $\phi_1, \ldots, \phi_n$ be an orthonormal basis for $\mathcal{F}$. Let $\alpha_i = \phi_i^{\mathcal{A}}$ for all i, and let $\varepsilon_i = \phi_i - \alpha_i$. We will call ϕ, α, and ε the *gauge*, *approximation*, and *error bases*, respectively (even though α_i and ε_i need not be independent). The average error map $\sigma_{p,\mathcal{A},\mathcal{F}}(x)$ can be expressed in terms of the error basis:

$$\begin{aligned}
\sigma_{p,\mathcal{A},\mathcal{F}}(x) &= \left[\int_{\mathbf{S}^{n-1}} \left|(\sum_i c_i\phi_i)(x) - (\sum_i c_i\phi_i)^{\mathcal{A}}(x)\right|^p dc\right]^{1/p} \Big/ \left[\int_{\mathbf{S}^{n-1}} 1\, dc\right]^{1/p} \\
&= \left[\frac{1}{A_n}\int_{\mathbf{S}^{n-1}} \left|\sum_i c_i\varepsilon_i(x)\right|^p dc\right]^{1/p}, \qquad (5)
\end{aligned}$$

where $A_n = 2\pi^{\frac{n}{2}}/\Gamma(\frac{n}{2})$ is the measure of $\mathbf{S}^{n-1}$.

Note that $\sum_i c_i\varepsilon_i(x)$ is the dot product of the unit vector $c = (c_1, c_2, \ldots, c_n)$ and the vector $\varepsilon(x) = (\varepsilon_1(x), \varepsilon_2(x), \ldots \varepsilon_n(x))$; it depends only on $|\varepsilon(x)|$ the angle θ between those two vectors, and is constant over the slice of $\mathbf{S}^{n-1}$ where θ is constant. The measure of that slice is $A_{n-1}\,|\sin\theta|^{n-1}\, d\theta$. Therefore,

$$\begin{aligned}
\sigma_{p,\mathcal{A},\mathcal{F}}(x) &= \left[\frac{1}{A_n}\int_0^\pi |\varepsilon(x)|\,|\cos\theta|^p\, A_{n-1}\,|\sin\theta|^{n-1}\, d\theta\right]^{1/p} \\
&= |\varepsilon(x)| \left[\frac{A_{n-1}}{A_n}\int_0^\pi |\cos\theta|^p\,|\sin\theta|^{n-1}\, d\theta\right]^{1/p} \\
&= |\varepsilon(x)| \left[\frac{(\Gamma(\frac{n}{2}))^2\Gamma(\frac{p+1}{2})}{\sqrt{\pi}\,\Gamma(\frac{n-1}{2})\Gamma(\frac{p+1+n}{2})}\right]^{1/p}. \qquad (6)
\end{aligned}$$

3.2. Explicit formula for μ. The worst-case error map $\mu_{\mathcal{A},\mathcal{F}}$ can be obtained by taking p to the limit $+\infty$ in formula (6), or directly, as follows. From formula (4),

$$\mu_{\mathcal{A},\mathcal{F}}(x) = \sup\left\{ \left|(\sum_i c_i\phi_i)(x) - (\sum_i c_i\phi_i)^{\mathcal{A}}(x)\right| \;:\; \left\|\sum_i c_i\phi_i\right\| = 1 \right\}$$

$$= \sup\Big\{ \Big|\sum_i c_i\varepsilon_i(x)\Big| \,:\, c \in \mathbf{S}^{n-1} \Big\}. \tag{7}$$

By considering the effect of negating each c_i, it is easy to see that the absolute value in the last formula is superfluous, i.e.

$$\mu_{\mathcal{A},\mathcal{F}}(x) \;=\; \sup\Big\{ \sum_i c_i\varepsilon_i(x) \,:\, c \in \mathbf{S}^{n-1} \Big\}. \tag{8}$$

Formula (8) is the supremum of a linear functional with coefficients $\varepsilon_i(x)$ over the sphere $\mathbf{S}^{n-1}$; which is achieved at the point $c^*(x)$ of $\mathbf{S}^{n-1}$ that is collinear with the coefficient vector, namely $c_i^*(x) = \varepsilon_i(x)/\sqrt{\sum_j(\varepsilon_j(x))^2}$, whence

$$\mu_{\mathcal{A},\mathcal{F}}(x) \;=\; \sum_i c_i^*(x)\varepsilon_i(x) \;=\; \sqrt{\sum_j (\varepsilon_j(x))^2} \;=\; |\varepsilon(x)|\,. \tag{9}$$

In summary, the error maps $\sigma_{p,\mathcal{A},\mathcal{F}}(x)$ and $\mu_{\mathcal{A},\mathcal{F}}(x)$ (which differ only by a constant factor) can be derived from the approximation errors $\varepsilon_i(x)$ for each basis function $\phi_i(x)$, combined with the norm $|\varepsilon(x)| = \sqrt{\sum_i(\varepsilon_i(x))^2}$.

4. Practical considerations

4.1. Connection between the function and point norms. The maps (4) and (3) will be more useful when there is a direct connection between the function-space norm $\|\cdot\|$ and the absolute value $|\cdot|$, used to compare functions values at a given point x, as in formulas (3)–(4)—namely, when

$$\|f\| = \left[\int_\Omega |f(x)|^q \; dx\right]^{1/q}. \tag{10}$$

More generally, the function values at x could be compared with a norm which could depend on x, or take derivatives of the function into account. We will not pursue such extensions in this paper.

Connection (10) is not strictly necessary—at least when $\mathcal{A}$ and $\mathcal{F}$ are finite dimensional. However, it may not make much sense to choose the approximant $f^{\mathcal{A}}$ so as to minimize the function norm $\|\cdot\|$, and then analyze its accuracy using some other norm $|\cdot|$, if there is no connection between the two.

Considering that the error map is relatively easy to compute when $\|\cdot\|$ is the L_2 norm (see section 3), and probably intractable otherwise, the connection expressed by formula (10) will probably hold in practice (with $q = 2$).

4.2. Choice of the gauge space. The approximation error map depends not only on the space $\mathcal{A}$, but also on the gauge space $\mathcal{F}$ and the error metric $\|f\|$. Therefore, the choice of $\mathcal{F}$ and $\|\cdot\|$ must be guided by the intended application.

For example, suppose the domain Ω is the circle or the sphere $\mathbf{S}^d$, and the application does not specify a preferred direction. Then we should choose $\mathcal{F}$ and $\|\cdot\|$ so that they are invariant under rotations of Ω—otherwise, any inhomogeneity in them may produce irrelevant artifacts in the error map. Also, if the functions to be approximated are expected to be smooth, and/or only their low frequencies are important, then the function in $\mathcal{F}$ should be smooth too. A natural choice for $\mathcal{F}$, in this case, are the circular or spherical harmonics up to a certain maximum order, and the metric $\|\cdot\|$ can be simply the L_q norm over the sphere $\mathbf{S}^d$.

4.3. Essential dimensions. We will argue next that, for the L_2 function norm, the "interesting" part of the error map is determined by two "essential" subspaces $\mathcal{F}' \subseteq \mathcal{F}$ and $\mathcal{A}' \subseteq \mathcal{A}$, which are disjoint and such that $\dim \mathcal{F}' \geq \dim \mathcal{A}'$.

First, if the spaces $\mathcal{A}$ and $\mathcal{F}$ have a non-trivial intersection $\mathcal{V}$, and we split a function $f \in \mathcal{F}$ into its components $g \in \mathcal{V}$ and $h \perp \mathcal{V}$, we find that $f^{\mathcal{A}} = g + h^{\mathcal{A}}$; and that $h^{\mathcal{A}}$ is itself orthogonal to $\mathcal{V}$. Therefore, we can confine our attention to the complements $\mathcal{F}'$ and $\mathcal{A}'$ of $\mathcal{V}$ relative to $\mathcal{A}$ and $\mathcal{F}$, which are disjoint.

Let us then suppose that $\mathcal{A}$ and $\mathcal{F}$ are disjoint. If $\dim \mathcal{F} < \dim \mathcal{A}$, let $\mathcal{A}' \subset \mathcal{A}$ be the projection of $\mathcal{F}$ onto $\mathcal{A}$, which contains all optimum approximants. Obviously, for any function f, we have $f^{\mathcal{A}} = f^{\mathcal{A}'}$, so we can confine our attention to the space $\mathcal{A}'$, which is still disjoint from $\mathcal{F}$ and satisfies $\dim \mathcal{F} \geq \dim \mathcal{A}'$.

5. Examples

5.1. Trigonometric splines on the circle. Consider the approximation of a function by continuous trigonometric splines, of maximum frequency $r = 2$, defined on a partition T of $\mathbf{S}^1$ into $n = 8$ *unequal* intervals. This space coincides with the space $\mathcal{P}_0^{r,2}[T]$ of non-homogeneous polynomial splines of $\mathbf{R}^2$, restricted to $\mathbf{S}^1$, with C_0 continuity constraints [2].

For the gauge space $\mathcal{F}$, we will use the family of trigonometric series truncated after a suitable maximum frequency $s \geq r$, which coincides with the space of general spherical polynomials (not splines) $\mathcal{P}^{s,2}$ for some $s \geq r$. The norm is $\|f\| = \sqrt{\langle f, f\rangle}$ where $\langle f, g\rangle = \int_{\mathbf{S}^1} f(\theta) g(\theta)\, dp$. Specifically, T consists of the intervals I_0 through I_7 shown below

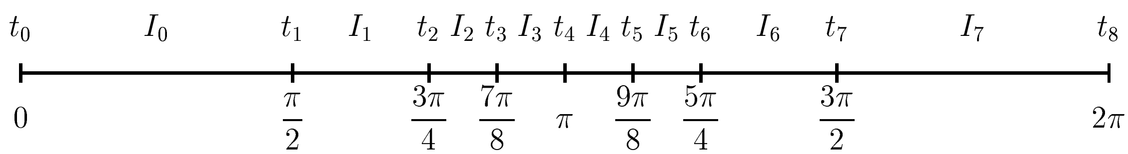

Within each interval I_j, the generic approximant is a linear combination g_j of the Fourier basis functions ϕ_i, for $-r \leq i \leq +r$. These partial functions are constrained to be continuous across interval boundaries; i.e. $g_{j-1}(t_j) = g_j(t_j)$ for each j in $\{0, \ldots, n-1\}$ (where all indices are taken modulo n). These equations turn out to be independent, therefore the dimension of $\mathcal{A}$ is $n(2r+1) - n = 32$.

For the gauge space $\mathcal{F}$, we will use the trigonometric polynomials of some order $s \geq r$, i.e. linear combinations of the basis functions ϕ_i for $-s \leq i \leq +s$, where $\phi_i(\theta) = (1/\sqrt{\pi})\sin(i\theta + \pi/4)$. As observed in section 4.3, we can ignore the subspace $\mathcal{A}' = \mathcal{F} \cap \mathcal{A}$ of $\mathcal{A}$ generated by $\phi_{-r}, \ldots, \phi_r$. Moreover, in order to use all of $\mathcal{A}$, we need $\dim \mathcal{F} \geq \dim \mathcal{A}$—i.e., $2s + 1 \geq 32$, implying $s \geq 16$. See figure 1. The resulting error map $\mu_{\mathcal{A},\mathcal{F}}(x)$ is shown in figure 2.

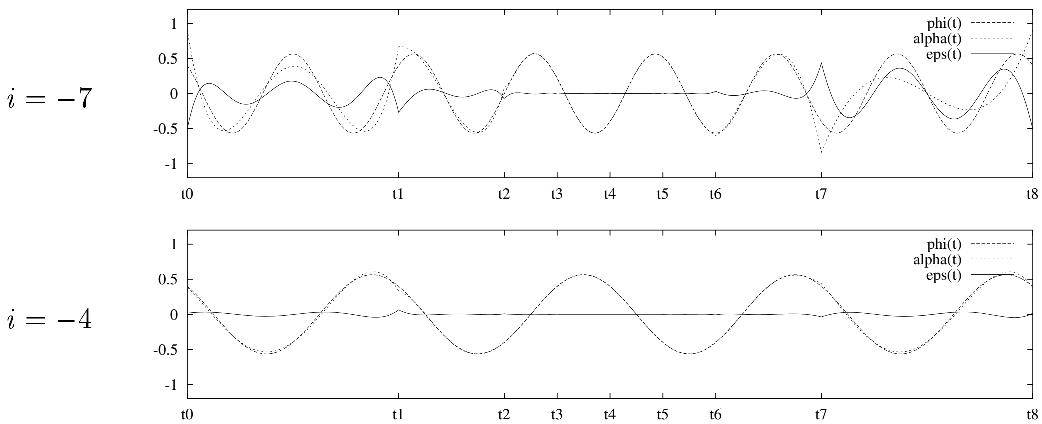

Figure 1: The functions $\phi_i(t)$, $\alpha_i(t)$, and $\varepsilon_i(t)$, for selected values of i.

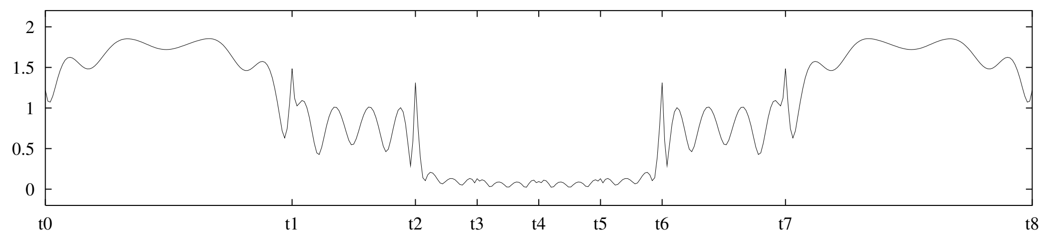

Figure 2: The error map $\mu_{\mathcal{A},\mathcal{F}}(t)$ for continuous (C_0) trigonometric splines on eight unequal intervals, tested with the space of trigonometric polynomials of order 16.

5.2. Spherical splines on a uniform mesh. For the examples in this section, the approximating functions are spherical polynomial splines [1, 2, 3, 4] of continuity class zero and various degrees, homogeneous and non-homogeneous, defined on some triangulation T of the sphere $\mathbf{S}^2$.

Figure 3 (left) shows the approximation error map $\mu_{\mathcal{A},\mathcal{F}}(p)$ for the *homogeneous* spherical spline space $\mathcal{A} = \mathcal{H}_0^5[T]/\mathbf{S}^2$, which has dimension 252. In figure 3 (right), $\mathcal{A}$ is the *non-homogeneous* spherical spline space $\mathcal{P}_0^4[T]/\mathbf{S}^2$, which has dimension 254. In both cases, the gauge space $\mathcal{F}$ is the family $\mathcal{Y}_{15}^2$ of spherical harmonics of maximum order 15, which has dimension 256. The intersection $\mathcal{F} \cap \mathcal{H}_0^5[T]/\mathbf{S}^2$ is the family of spherical harmonics of odd order ≤ 5 (dimension 21), whereas $\mathcal{F} \cap \mathcal{P}_0^4[T]/\mathbf{S}^2$ is the full harmonic space $\mathcal{Y}_4^2$ (dimension 25). The level curves are logarithmically spaced, five per decade.

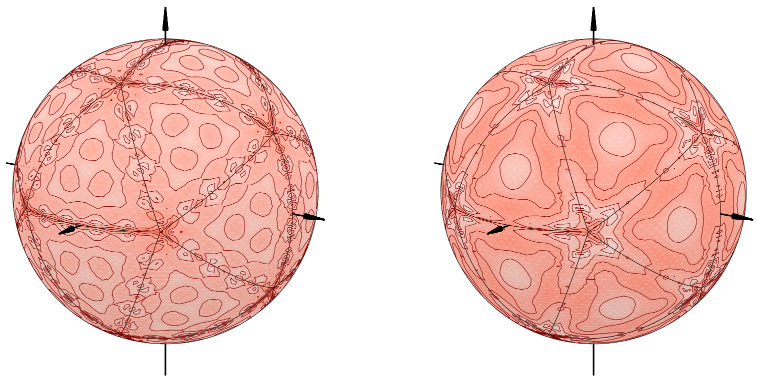

Figure 3: Error maps $\mu_{\mathcal{A},\mathcal{F}}(p)$ for the approximation spaces $\mathcal{A} = \mathcal{H}_0^5[T]/\mathbf{S}^2$ (left) and $\mathcal{A} = \mathcal{P}_0^4[T]/\mathbf{S}^2$ (right). The maximum errors are 13.5 and 9.37, respectively.

5.3. Spherical splines on a variable mesh. In the following examples, the approximating functions are again spherical polynomial splines, but the vertices of the triangulation T have been displaced so as to create regions of very different sizes (still with icosahedral topology).

Figure 4 (left) shows the approximation error map $\mu_{\mathcal{A},\mathcal{F}}(p)$ for the space of *homogeneous* spherical splines $\mathcal{A} = \mathcal{H}_0^5[T]/\mathbf{S}^2$, which has dimension 252. In figure 4 (right), $\mathcal{A}$ is the space of *non-homogeneous* spherical splines $\mathcal{P}_0^4[T']/\mathbf{S}^2$, which has dimension 254. In both cases, the gauge space $\mathcal{F}$ is the family $\mathcal{Y}_{15}^2$ of spherical harmonics of maximum order 15, which has dimension 256, as before. The level curves are logarithmically spaced (5 per decade).

6. Conclusion

Asymptotic error analysis is not very helpful when comparing two fixed finite-dimensional approximation spaces of similar dimensions—such as a spline space against a wavelet space, or two spline spaces with different grid geometries. Approximation errors computed for individual test functions are difficult to interpret and may not be representative of the average or worst cases. We expect that the approximation error map will be a useful analysis tool for those situations—especially for domains that admit natural uniform target spaces, such as spheres (including the circle) and tori.

Acknowledgments. This research was supported in part by CAPES, FINEP, and CNPq (PRONEX-SAI).

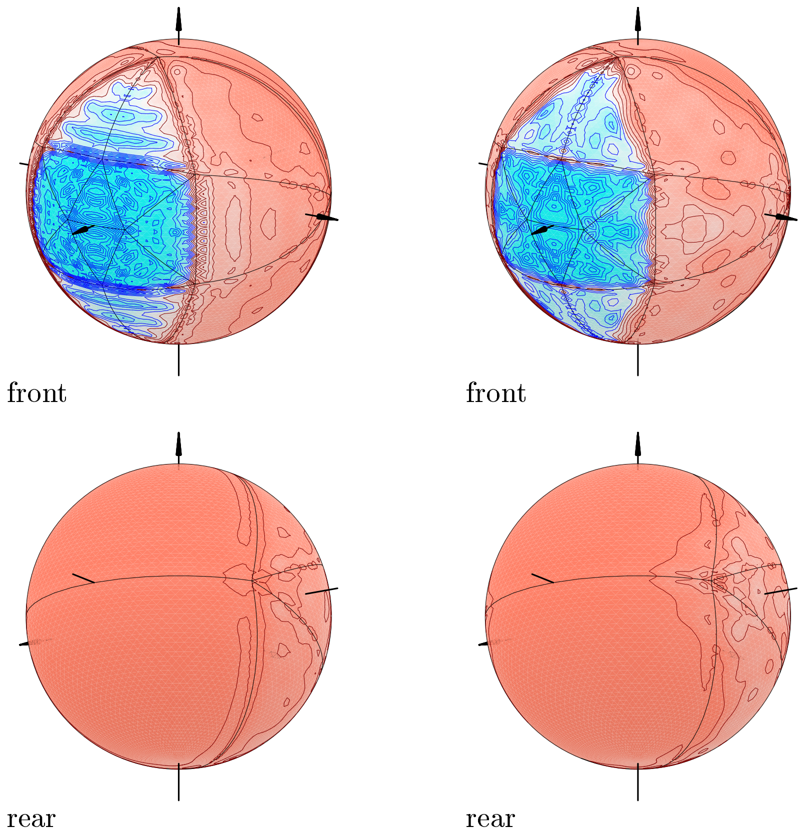

Figure 4: Error maps $\mu_{\mathcal{A},\mathcal{F}}(p)$ for the approximation spaces $\mathcal{A} = \mathcal{H}_0^5[T]/\mathbf{S}^2$ (left) and $\mathcal{A} = \mathcal{P}_0^4[T]/\mathbf{S}^2$ (right). The maximum errors are 17.1 and 17.9, respectively.

References

[1] P. Alfeld, M. Neamtu, and L. L. Schumaker. Dimension and local bases of homogeneous spline spaces. *SIAM Journal of Mathematical Analysis*, 27(5):1482–1501, Sept. 1996.

[2] A. Gomide and J. Stolfi. Non-homogeneous polynomial C_k splines on the sphere S^n. Technical Report IC-00-10, Institute of Computing, Univ. of Campinas, July 2000.

[3] A. Gomide and J. Stolfi. Bases for non-homogeneous polynomial $\mathbf{C}_k$ splines on the sphere. In *Lecture Notes in Computer Science 1380: Proc. LATIN'98 - Latin American Theoretical Informatics Conference*, pages 133–140. Springer, Apr. 1998.

[4] A. Gomide. *Splines Polinomiais Não Homogêneos na Esfera*. PhD thesis, Institute of Computing, University of Campinas, May 1999. (In Portuguese).