Noname manuscript No.
(will be inserted by the editor)

Numerical approximation of the fractional Laplacian via *hp*-finite elements, with an application to image denoising

Paolo Gatto · Jan S. Hesthaven

the date of receipt and acceptance should be inserted later

Abstract The fractional Laplacian operator $(-\Delta)^s$ on a bounded domain Ω can be realized as a Dirichlet-to-Neumann map for a degenerate elliptic equation posed in the semi-infinite cylinder $\Omega \times (0, \infty)$. In fact, the Neumann trace on Ω involves a Muckenhoupt weight that, according to the fractional exponent s, either vanishes ($s < 1/2$) or blows up ($s > 1/2$). On the other hand, the normal trace of the solution has the reverse behavior, thus making the Neumann trace analytically well-defined. Nevertheless, the solution develops an increasingly sharp boundary layer in the vicinity of Ω as s decreases. In this work, we extend the technology of automatic *hp*-adaptivity, originally developed for standard elliptic equations, to the energy setting of a Sobolev space with a Muckenhoupt weight, in order to accommodate for the problem of interest. The numerical evidence confirms that the method maintain exponential convergence. Finally, we discuss image denoising via the fractional Laplacian. In the image processing community, the standard way to apply the fractional Laplacian to a corrupted image is as a filter in Fourier space. This construction is inherently affected by the Gibbs phenomenon, which prevents the direct application to "spliced" images. Since our numerical approximation relies instead on the extension problem, it allows for processing different portions of a noisy image independently and combine them, without complications induced by the Gibbs phenomenon.

Keywords fractional Laplacian · *hp*-finite elements · automatic adaptivity · image denoising

Mathematics Subject Classification (2000) 65N30 · 35J75

P. Gatto
Division of Applied Mathematics, Brown University, 182 George Street, Providence, RI 02912, U.S.A.
E-mail: paolo_gatto@brown.edu

J.S. Hesthaven
Ecole Polytechnique Fédérale de Lausanne (EPFL) CH-1015 Lausanne, Switzerland
E-mail: jan.hesthaven@epfl.ch

1 Introduction

Fractional PDE's have recently become a topic of substantial research activity because of their ability to describe phenomena with memory effects which elude the reach of standard PDE's. Among others, fractional calculus has proved a viable tool in material science [2], finance [7], image processing [3], porous media flow [22], and bioengineering applications, such as anomalous diffusion in biological tissues [16].

The fractional Laplacian operator is of special interest in fractional calculus. Caffarelli and Silvestre, [5], have shown that such operator can be realized as a Dirichlet-to-Neumann map for a degenerate elliptic equation posed on the half-space $\mathbb{R}^+_n := \mathbb{R}^n \times (0, \infty)$. The degeneracy is caused by a weight function that, depending upon the fractional exponent, either vanishes or blows up on the boundary $\partial\mathbb{R}^+_n = \mathbb{R}^n$. This result has been adapted to the case of a bounded domain $\Omega \subset \mathbb{R}^n$ in [4], [6], [21], thus obtaining a degenerate elliptic boundary value problem problem posed on the infinite cylinder $\mathcal{C} := \Omega \times (0, \infty)$. Since the weight function belongs to the so-called Muckenhoupt class, the problem can be understood in the framework of weighted Sobolev spaces, and the standard theory carries through with minor modifications.

By employing separation of variables, see Appendix A, it can be shown that the solution to the extension problem decays to zero away from Ω, thus allowing for the truncation of $\mathcal{C}$ to a (large) bounded domain, and the approximation of the problem through numerical techniques such as *hp*-finite elements. At the same time, due to the fractional exponent, either the Muckenhoupt weight or the gradient of the solution blows up on Ω. The flexibility of varying both h and p, along with an automatic adaptive strategy, could ensure that the solution is properly captured in a neighborhood of Ω.

The paper is organized as follows. In Section 2 we recall theoretical results about the fractional Laplacian, including a trace theorem which was proved by Nochetto et al. [18]. Section 3 is devoted to numerical experiments that were conducted using the `hp2d` finite element library [8]. In Section 4 our numerical machinery is employed for the solution of the space-fractional diffusion equation, which is used as a tool for image denoising. Finally, in Section 5 we draw conclusions from this work.

2 Preliminary Results

The fractional Laplacian $(-\Delta)^s$ is a pseudo-differential operator defined by

$$\mathscr{F}[(-\Delta)^s u](\xi) = |\xi|^{2s} \mathscr{F}[u](\xi) \tag{1}$$

where $\mathscr{F}[\cdot]$ indicates the Fourier transform. When u belongs to the Schwartz class $\mathcal{S}$ of functions on $\mathbb{R}^n$ and $s > -n/2$, i.e., the right-hand-side of (1) is a tempered distribution, and $(-\Delta)^s u$ is a tempered distribution as well. By duality, the previous definition extends to a subspace of the space of tempered distributions $\mathcal{S}'$, see, e.g., Silvestre [19]. In the case $-n/2 < s < 0$, [20], the fractional Laplacian can be expressed as the convolution integral with the Riesz kernel, i.e.,

$$(-\Delta)^s u(x) = C_{n,s} \int_{\mathbb{R}^n} \frac{u(y)}{|x-y|^{n-2s}}\, dy \tag{2}$$

where $C_{n,s}$ is a suitable constant.

When $\mathbb{R}^n$ is replaced by a bounded domain Ω, there is no unique way to define the fractional Laplacian. For example, a possible definition that rests upon restricting the integral in (2) to Ω, would lead to the so-called "regional" fractional Laplacian [13]. In this work, we pursue a different approach and consider the spectral decomposition of the homogeneous Dirichlet Laplace operator to define fractional powers of such an operator through classical spectral theory, [4,6]. When Ω is a Lipschitz domain, as we shall henceforth assume, the spectrum of $-\Delta$ is discrete and positive, and accumulates at infinity. Thus, we formally define

$$(-\Delta)^s u = \sum_k \mu_k^s u_k \varphi_k \tag{3}$$

where $u_k = \int_\Omega u\varphi_k$ are the Fourier coefficients of u, and $\{\mu_k\}$, $\{\varphi_k\}$ are the eigenvalues and eigenvectors of $-\Delta$, respectively.

Since $-\Delta : H_0^1(\Omega) \to L^2(\Omega)$ is an unbounded self-adjoint operator, by virtue of the theory of Hilbert scales presented by Magenes and Lions [15], we can define the interpolation space

$$\left[H_0^1(\Omega), L^2(\Omega)\right]_\theta := \operatorname{dom}\left((-\Delta)^{\frac{1-\theta}{2}}\right) \tag{4}$$

for $0 \le \theta \le 1$. Furthermore, the space on the left-hand side of (4) is the space $\widetilde{H}^s(\Omega)$, where $s = 1-\theta$. Let us recall that, if $H^s(\mathbb{R}^n)$ is the scale of Sobolev spaces defined through the Bessel potential (see [17], Chapter 3) and $H^s(\Omega)$ is the space of restrictions to an arbitrary domain Ω, we define

$$\begin{aligned} H_0^s(\Omega) &\qquad \text{closure of } \mathscr{D}(\Omega) \text{ in } H^s(\Omega) \\ \widetilde{H}^s(\Omega) &\qquad \text{closure of } \mathscr{D}(\Omega) \text{ in } H^s(\mathbb{R}^n) \end{aligned}$$

The definition trivially implies that $\widetilde{H}^s(\Omega) \subset H_0^s(\Omega)$ and, in the case of Ω being a Lipschitz domain and s being different from a half-integer, it can be shown that the two spaces coincides.

The fractional Laplacian, see [4,6,21], can be realized as the restriction to Ω of the following degenerate elliptic problem on the semi-infinite cylinder $\mathcal{C} := \Omega \times (0,\infty)$:

$$-\operatorname{div}(y^\alpha \nabla u) = 0 \qquad \text{in } \Omega \times (0,\infty) \tag{5a}$$

$$u = 0 \qquad \text{on } \partial\Omega \times [0,\infty) \tag{5b}$$

$$(y^\alpha \nabla u)\cdot \boldsymbol{\nu} = d_s f \qquad \text{on } \Omega \times \{0\} \tag{5c}$$

Here $\alpha = 1-2s$, $y \in \mathbb{R}^+$ is the extension variable, $\boldsymbol{\nu}$ is the external normal vector to $\Omega \times (0,\infty)$, $d_s = \Gamma(1-s)\frac{1}{2}^{s+1}$ is a constant, and Eqs (5b) (5c) are to be understood in the sense of the trace in a weighted Sobolev space. Then, for $f \in L^2(\Omega)$, it is possible to show that $u|_\Omega$ solves $(-\Delta)^s u = f$ in Ω and $u = 0$[1] on $\partial\Omega$, see Appendix A for a detailed discussion.

Problem (5) is degenerate in the sense that the coefficient y^α does not belong to $L^\infty(\mathcal{C})$ as in the standard theory. On the other hand, since the function $w := y^\alpha$ belongs to the Muckenhoupt class A_2, we can recast the problem in the weighted Sobolev space $H^1(\mathcal{C}; w)$, endowed with the inner product

$$(u,v) = \int_{\mathcal{C}} w\,\nabla u \cdot \nabla v + \int_{\mathcal{C}} w\,uv$$

[1] Additional regularity on u is in fact needed for this condition to be well-defined. Because of a shift-type result [14], it is sufficient to assume $f \in H^t(\Omega)$, for $t > 1/2 - 2s$.

and the corresponding norm $\|u\|^2 = (u,u)$. For a theory of Sobolev spaces with Muckenhoupt weights, see [12]. Using integration by parts and proceeding formally, (5) is equivalent to the following variational problem:

$$\text{Find } u \in H^1_{0,\Gamma_0}(\mathcal{C};w) \;\; : \;\; \int_{\mathcal{C}} y^\alpha \, \nabla u \cdot \nabla v = d_s \int_\Omega f v \qquad \forall \, v \in H^1_{0,\Gamma_0}(\mathcal{C};w) \tag{6}$$

where $H^1_{0,\Gamma_0}(\mathcal{C};w)$ is the space of objects in $H^1(\mathcal{C};w)$ with null trace on $\Gamma_0 := \partial\Omega \times (0,\infty)$. Since w does not depend on $x \in \Omega$, a trace result on Γ_0 trivially holds. In the case of the trace on Ω, it was shown in [18] that $v \mapsto v|_\Omega$ is a continuous operator of $H^1(\mathcal{C},w)$ onto $\widetilde{H}^s(\Omega)$. Therefore the right-hand side of the variational equality is well defined for every load f that belongs to $H^{-s}(\Omega) = \widetilde{H}^s(\Omega)'$, and the integral is indeed interpreted as a duality pairing. Finally, since a Poincaré-type inequality for $H^1_{0,\Gamma_0}(\mathcal{C};w)$ holds, [18], the Lax-Milgram Lemma implies the existence of a unique solution to (6).

3 Finite Elements Approximation

The variational formulation (6) immediately allows for a finite element approximation of the fractional Laplacian. The asymptotic expansion of the Bessel function of the second kind K_s (Appendix A) implies that the solution develops a boundary layer as we approach Ω. On the other hand, as shown in [18], ∇u decays exponentially as $y \uparrow \infty$. Thus, we can select a sufficiently large value of y, replace $\mathcal{C}$ by a bounded computational domain $\mathcal{C}_{hp} = \Omega \times (0,L)$, and impose a homogeneous Dirichlet condition on $\Omega \times \{y = L\}$. Alternatively, we could employ an Absorbing Boundary Condition to terminate the computational domain.

We shall rely upon the `hp2d` finite element library [10], to conduct numerical experiments. This choice is motivated by the fact that `hp2d` supports automatic *hp*-adaptivity for elliptic problems, and optimal *h*-refinements and *p*-enrichments are automatically selected via an error indicator that relies upon H^1-projection based interpolation, see [9].

For the purpose of numerical simulations, we choose $\Omega = (0,1)$, $\mathcal{C}_{hp} = (0,1)^2$, and employ the manufactured solution

$$u = \begin{cases} \sin(\pi x)\, \exp(-\pi y) & s = \frac{1}{2} \\ \sin(\pi x)\, (\pi y)^s \, K_s(\pi y) & \text{otherwise} \end{cases}$$

for computing the load function $f = \frac{1}{d_s} \lim_{y \downarrow 0} y^\alpha u_y$. The Bessel function of the second kind K_s is evaluated using the implementation of Amos, see [1]. Although we employ the exact solution to impose an non-homogenous boundary condition on $\Omega \times \{y = 1\}$, numerical evidence suggests that we could assume a homogeneous condition as well, thus validating our strategy of domain truncation.

Let u_0 be the restriction of the exact solution u on $\Omega \times \{y = 1\}$, let V_h be an H^1-conforming finite element space on $\mathcal{C}_{hp}$ and let $u_{0,h}$ be a lift of u_0 in V_h. The finite element approximation of (6) is:

$$\text{Find } u_h \in u_{0,h} + V_h \;\; : \;\; \int_{\mathcal{C}_{hp}} y^\alpha \, \nabla u_h \cdot \nabla v_h = d_s \int_\Omega f \, v_h \qquad \forall \, v_h \in V \tag{7}$$

We employ adaptive quadrature for the integration of both the stiffness matrix and the load vector. The `hp2d` library supports fully variable order elements. This means that a quadrilateral element has independent orders of approximation for each edge and element interior which, because of the tensor

product structure, has a horizontal and vertical order of approximation. The initial mesh is comprised of just one isotropic second order quadrilateral element.

The adaptive strategy of the `hp2d` library is based on the idea of comparing the approximate solution u_{hp} on the current grid to a solution $u_{h/2\,p+1}$ on a finer grid in order to determine an optimal mesh for the next step. The optimal mesh is understood as a mesh that maximizes the rate of decrease of a properly constructed error indicator. The existing implementation was expanded in order to account for the appropriate energy setting. Ideally, the *hp*-algorithm terminates when a given tolerance ε, defined as:

$$\frac{\left|u_{hp} - u_{h/2\,p+1}\right|_{w,1}}{\left|u_{h/2\,p+1}\right|_{w,1}} < \varepsilon$$

is reached. Here $|u|^2_{w,1} = \int w\,|\nabla u|^2$ indicates the seminorm in $H^1_{0,\Gamma_0}(\mathcal{C}_{hp}; w)$.

The objective of the numerical experiments is to study the convergence rates of $\left|u-u_h\right|$ and $\left|(u-u_h)|_\Omega\right|$ in $H^1_{0,\Gamma_0}(\mathcal{C}_{hp}; w)$ and $\widetilde{H}^s(\Omega)$, respectively. According to the trace theorem from Section 2, these rates shall coincide. In the context of *h*-estimates for $\mathbb{Q}^1$ finite element approximations, the following result has been proved in [18]:

$$|u - u_h|_{w,1} \leq O\Big(\#\,\text{degrees of freedom}^{-\frac{1}{n+1}}\Big)$$

where n is the Euclidean dimension of Ω. This estimate was obtained for graded meshes, hence the use of the number of degrees of freedom (dof's) rather than the mesh size h. To the best of our knowledge, no *p*-estimate is available, thus no *hp*-estimate can be immediately obtained through a standard scaling argument. The numerical results suggest that exponential convergence is achieved by our numerical scheme.

The case $s = 1/2$, Figure 1, reduces to the standard Laplacian with a Neumann boundary condition on Ω and will be regarded as a benchmark for the fractional Laplacian. The tolerance $\varepsilon = 0.01$ is reached within 4 iterations of the *hp*-algorithm on a mesh with 30 dof's, see Figure 1(a). Because the solution to the problem is regular, the *hp*-algorithm has produced only *p*-enrichments. Figure 1(b) illustrates the convergence curves, obtained by forcing 25 iterations. The convergence curves for the error and error indicator are practically indistinguishable. The trace of the error on Ω exhibits the same behavior as the error and error indicator on $\mathcal{C}_{hp}$.

When $s < 1/2$, causing the gradient of the solution u on Ω to become increasingly singular as s approaches 0, we shall expect the method to degrade as smaller and smaller values of s are selected. We set $\varepsilon = 0.01$ and perform numerical experiments for $s = 0.48$, 0.25, and 0.05, illustrated in Figures 2–4. For $s = 0.48$, the desired tolerance is reached within 7 iterations of the *hp*-algorithm on a mesh with 165 dof's, see Figure 2(a). In the case $s = 0.25$, the algorithm needs 27 iterations to reach the desired tolerance on a mesh with 5817 dof's, see Figure 3(a). The final mesh features heavy *h*-refinements in a layer adjacent to Ω, while high orders of approximation are employed in the elements farther away from Ω. In both cases, see Figure 2(b) and Figure 3(b), the convergence appears to be of exponential type. In the case $s = 0.25$, because of the nature of adaptive quadrature, the heavy *h*-refinements that occur where the gradient of the solution is singular cause the computation of the error to eventually fail. This problem could be circumvented by limiting adaptive quadrature in a neighborhood of Ω. Nevertheless, we are confident that the results are meaningful because the error and error indicator exhibit the same behavior for the first 21 iterations. Lastly, in the case $s = 0.05$, after 40 iterations the *hp*-algorithm terminates at $\varepsilon = 0.1089$ on a mesh with 13576 dof's, see Figure 4(a). Despite failing to reach the desired error tolerance, the trace of the error on Ω reaches

a relative error of less than 1% in 27 iterations, on a mesh with 3959 dof's. As in the previous case, the computation of the norm of the error eventually fails. Although exponential convergence on the extended domain is lost, the trace of the error on Ω seems to attain nearly exponential convergence, see Figure 4(b).

In the case $s > 1/2$, because of the increasingly singular behavior of the weight function y^{1-2s} as s approaches 1, we expect a deterioration of the numerical results as s increases. We select $\varepsilon = 0.01$ as before, and perform numerical experiments for $s = 0.55$, 0.65, and 0.75, see Figures 5–7. In the first case, 12 iterations and 296 dof's are needed to reach the desired tolerance. Because of the smoothness of the solution, the *hp*-algorithm has favored *p*-enrichments, see Figure 5(a). To highlight convergence rates, 25 iterations of the *hp*-algorithm are performed. Figure 5(b) suggests that, after a pre-asymptotic region, the error, the error indicator, and the trace of the error on Ω exhibit exponential convergence. When $s = 0.65$, the desired tolerance is reached within 16 iteration on a mesh with 774 dof's, see Figure 5(a). The behavior of the error, the error indicator and the trace of the error on Ω is similar to the previous case, see Figure 5(b) for results over 25 iterations. In the last case, 15 iterations of the *hp*-algorithm and a final mesh with 451 dof's, see Figure 7(a), are needed to reach the desired tolerance. We conjecture that the decrease in the number of dof's is a consequence of the increased regularity of the gradient of the solution, which overcompensates for the singular behavior of the weight function. Figure 7(b) shows the behavior of the error, the error indicator and the trace of the error on Ω for 25 iterations. Although good agreement between the error and error indicator is observed, as a consequence of the increased blow up of the weight function, the trace of the error on Ω begins to drift away. For completeness, we ran a convergence test for a larger value of s, namely $s = 0.85$, see Figure 8. While good agreement between the error and error indicator is still observed over the first few iterations, the behavior of the trace of the error on Ω becomes increasingly erratic because of the poor performance of adaptive integration. Let us emphasize that is not a limitation of the method, but rather a numerical integration issue that could be resolved with an appropriate quadrature strategy for singularities of the type $y^{-1+\epsilon}$.

4 Fractional Diffusion and Image Denoising

The goal of image denoising is to enhance the visual quality of an image by eliminating features which are random and uncorrelated. In general, noise tends to corrupt the higher frequency content of most images. A standard technique is to apply a suitable lowpass filter, e.g., a Gaussian function, to the image FFT to attenuate the high frequency components. However, because of the lack of periodicity, the final image exhibits a Gibbs phenomenon at its edges. Although this issue can be overcome by applying the filter to a periodic extension of the image, this same strategy does not apply to image splicing, and the Gibbs phenomenon will pollute the final image. Although techniques for removing the Gibbs phenomenon are known, [11], they are complicated and lack robustness.

A different approach, overcoming the Gibbs phenomenon altogether, is to apply the diffusion equation to the noisy image. By recalling the expression of the Green function for the heat operator, this is equivalent to apply a time-varying Gaussian filter to the noisy image. It is well-established in the literature that the heat operator diffuses too rapidly to preserve sharp edges. Thus, it is natural to consider the fractional diffusion equation:

$$\begin{aligned} \frac{\partial u}{\partial t} &= \Delta^s u \qquad \text{in } \Omega \times \{t > 0\} \\ u(x,0) &= u_0(x) \qquad \text{on } \Omega \times \{t = 0\} \end{aligned}$$

for some exponent $s < 1$. This approach has been followed in [3] in which the authors exploit the factorization $\Delta^s = \Delta\,\Delta^{s-1}$, and approximate the fractional part through filtering of the FFT, and the Laplacian through a standard finite difference scheme. We approximate Δ^s directly, via the extension problem.

A gray-scale image is a discrete function I that maps each pixel, identified to an element of the grid $\{1,\dots,m\}\times\{1,\dots,n\}$, to a value in the interval $[0,1]$. We employ our $\mathbb{R}\times(0,\infty)$ implementation by addressing each image slice $I_i := I_{i,1\to n}$ independently. For each slice I_i, we build the initial condition u_0 and approximate the fractional diffusion equation with an implicit time scheme, namely:

$$\frac{u^{k+1}-u^k}{\delta t} = \Delta^s u^{k+1} \qquad , \qquad k = 0,1,\dots$$

for a suitable time-step δt and exponent s. We shall exploit the possibility of varying both s and δt with each time-step iteration. Ideally, an optimal computational grid should be determined through the *hp*-adaptive scheme at each time step. In practice, we select a single fractional exponent $s = \overline{s}$ and produce one optimal grid to employ for each s. This choice is justified since s will vary over a modest range.

The extension problem requires the initial condition u_0 to vanish on $\partial\Omega$. We construct u_0 by subtracting the piecewise linear interplant defined by the knot vector

$$\Big\{(0,0)\,;\,(I_{i,1},\tfrac{1}{n+1})\,;\,(I_{i,n},\tfrac{n}{n+1})\,;\,(0,1)\Big\}$$

from the piecewise linear interpolant of I_i:

$$\Big\{(0,0)\,;\,(I_{i,1},\tfrac{1}{n+1})\,;\,\dots;(I_{i,n},\tfrac{n}{n+1})\,;\,(0,1)\Big\}$$

Thus, u_0 is a continuous function on $[0,1]$, vanishes on $[0,\frac{1}{1+n}]\cup[\frac{n}{n+1},1]$ and has the desirable feature of not having a sharp boundary layer in a neighborhood of the endpoints.

Appropriate values for s and δt are selected through an indicator E that measures the relative change in u between subsequent time-steps as compared to the relative change in the gradient of u. A natural definition for E is:

$$E(s,\delta t) = \frac{\big|\|u_k\|-\|u_{k+1}\|\big|}{\|u_k\|}\left(\frac{\|\nabla u_k\|}{\big|\|\nabla u_k\|-\|\nabla u_{k+1}\|\big|}\right)^2 \tag{8}$$

where $\|\cdot\|$ is a suitable norm. The choice of raising the inverse relative error in the gradient to the second power is arbitrary. In fact, any power strictly greater than one appears to be a reasonable choice. When $\|\cdot\|$ is the L^1 norm, then $E(s,\cdot)$ is the family of curves in Figure 9. The following strategy suggests itself: for each decreasing value of s, select δt that minimizes $E(s,\cdot)$. However, as s decreases, the "convexity" of each curve decreases, which makes finding its minimum increasingly challenging. Numerical experiments have shown that redefining the indicator as $E\,\delta t^s$ yields appropriate values for the time-step.

The adaptive denoising strategy is described in full details in Algorithm 1. It relies upon control parameters s_0, δt_0, ε_0, ε_1, ε_2 which are to be set manually. Pathological cases, such as a spike function and a step function, were employed to assess the performance of the strategy. The quality of an image is a subjective concept and the results should literally be evaluated according to the "eyeball" norm. In order to facilitate this, gray-scale images, whose vertical cross sections coincide with the original, corrupted and denoised functions were constructed as well. The results about the denoising of a

spike function are shown in Figure 10. While an improvement at each step is visible, the decrease of the relative L^1 error is only minimal, thus providing empirical evidence of the unique nature of the "eyeball" norm. Finally, a comparison between denoising via fractional and standard Laplacian in the case of a step function is made in Figure 11.

We applied the denoising algorithm to two images, see Figure 12 and Figure 13, and compared it to a similar algorithm that relies on the standard Laplacian. In the case of standard diffusion, based on numerical experiments, we chose $\|\cdot\|$ to be the L^2 norm. Because of this different energy setting, the correction factor δt^s turned out to be inappropriate. Each image was corrupted by adding a random noise uniformly distributed on $[-0.1, 0.1]$. As a first test, we picked the "peppers" image that is ubiquitous in the literature, see Figure 12; a second test was performed on an image of the David by Michelangelo, see Figure 13. The peppers image was manipulated by cutting it vertically into four strips, independently denoising each strip, and splicing it back together. In the case of the David, only three vertical strips were employed. As it can be observed from the images, the fractional Laplacian in fact preserves sharp edges better than the standard Laplacian, hence enhancing the contrast. Furthermore, no effects are present due to the vertical cuts where the splicing occurred. Finally, through a partition of unity, our methodology extends to the case of denoising images that partially overlap onto each other, which is indeed the situation encountered in several practical applications.

Algorithm 1: adaptive denoising strategy.

User provided parameters: $s_0, \delta t_0, \varepsilon_0, \varepsilon_1, \varepsilon_2$

construct u_0 from slice I_i;
set $s = s_0$, $\delta t = \delta t_0$;
solve for u_1;
set $s = (1 - \varepsilon_0)s$, $k = 1$, $E = huge$;
while $\|u_k\| \geq (1 - \varepsilon_1)\|u_1\|$ **do**
    solve for u_{k+1};
    compute E;
    if *E has decreased* **then**
        store u_{k+1};
        update E;
        $\delta t = (1 + \varepsilon_2)\delta t$;
    else
        retrieve previously determined u_{k+1};
        update $u_k = u_{k+1}$;
        set $E = huge$, $s = (1 - \varepsilon_0)s$, $k = k + 1$;
    end
end

5 Conclusions

The fractional Laplacian can be realized as a Dirichlet-to-Neumann map for a degenerate elliptic equation. The main objective of this work is to employ the technology of *hp*-adaptivity, developed

for standard elliptic problems, to obtain a robust numerical scheme for the solution of the fractional Laplacian in a bounded domain Ω. This was achieved by extending the *hp*-algorithm developed in [10,8], to the context of a Sobolev space with a weight in the Muckenhoupt class. Since the weight degenerates on a portion of the boundary, particular care had to be taken when evaluating the energy norms that drive the adaptive strategy. As illustrated through a number of numerical examples, the adaptive strategy produces meshes that exhibit, in a layer adjacent to the domain Ω, a fine grading that would otherwise be extremely challenging to determine with a manual, *ad hoc* strategy. Furthermore, although no theoretical result is available at the moment, we showed numerical evidence that *hp*-finite elements exhibit exponential convergence for this particular degenerate elliptic problem. Finally, as an application, we employed the fractional Laplacian for image denoising by solving a space-fractional diffusion equation. The implicit Euler method was used for performing the time-stepping. The results confirm the conjecture that the fractional Laplacian preserves sharper edges during the denoising process as compared to a standard diffusion equation. In the context of image processing, the fractional Laplacian is consistently understood and implemented as a filter in the frequency space, which makes it susceptible to the Gibbs phenomenon. Since our implementation relies instead on the extension problem, the Gibbs phenomenon is circumvented. Thus, our numerical approximation of the fractional Laplacian is suitable for image splicing.

Acknowledgements This work was partially supported by NSF DMS-1115416, by OSD/AFOSR FA9550-09-1-0613 and by AFOSR FA9550-12-1-0463. The authors would like to thank Johnny Guzman for many fruitful discussions.

A Extension problem: separation of variables

Our goal is to construct a solution to the extension problem (5) via separation of variables and show that the restriction of such a solution to Ω coincides with the solution of the fractional Laplacian as defined in (3) through the eigenmodes of the Dirichlet Laplace operator. Let us recall 5:

$$-\operatorname{div}(y^\alpha \nabla u) = 0 \qquad \text{in } \Omega \times (0,\infty) \tag{9a}$$

$$u = 0 \qquad \text{on } \partial\Omega \times [0,\infty) \ , \tag{9b}$$

$$(y^\alpha \nabla u)\cdot \boldsymbol{\nu} = d_s f \qquad \text{on } \Omega \times \{0\} \tag{9c}$$

where d_s is a suitable constant to be determined, and (9b) (9b) are understood in the sense of the trace. We shall assume $f \in L^2(\Omega)$. If we let $u(x,y) = X(x)Y(y)$, then equation (9a) can be separated as follows:

$$\frac{\alpha}{y}\frac{Y'}{Y} + \frac{Y''}{Y} = -\frac{\Delta X}{X} = c^2 \ .$$

Using the boundary condition (9b), it is immediate that X coincides with the eigenmodes $\{\varphi_k\}$ of the Dirichlet Laplace problem on Ω, and the separation constant c^2 must be an eigenvalue μ_k of such a problem. Thus, $X = \varphi_k$, and $c^2 = \mu_k$. With the separation constant determined, we move to the ordinary differential equation for Y. In the case $\alpha = 0$, i.e., $s = 1/2$, we obtain $Y = \exp(-\mu_k^{1/2} y)$. Namely, we discarded the solutions associated with $\exp(\mu_k^{1/2} y)$ because of the energy assumption $u \in H^1(\mathcal{C}, w)$. In the case $\alpha \neq 0$, function Y satisfies the following equation:

$$y^2\, Y'' + \alpha y\, Y' - (\mu_k^{1/2} y)^2 Y = 0 \ .$$

By seeking a solution Y of the form $(\mu_k^{1/2} y)^s\, g(\mu_k^{1/2} y)$, after lengthy yet trivial computations, we arrive at a modified Bessel equation for g:

$$(\mu_k^{1/2} y)^2 g''(\mu_k^{1/2} y) + (\mu_k^{1/2} y) g'(\mu_k^{1/2} y) - \big((\mu_k^{1/2} y)^2 + s^2\big) g(\mu_k^{1/2} y) = 0 \ .$$

The general solution is given by a linear combination of I_s and K_s, i.e., the modified Bessel function of the first kind and second kind, respectively. As in the previous case, because of the energy assumption, we discard the solutions associated to I_s. Thus, if we define:

$$\psi_k = \exp(-\mu_k^{1/2} y) \quad \text{for } s = \tfrac{1}{2} \qquad ; \qquad \psi_k = (\mu_k^{1/2} y)^s \, K_s(\mu_k^{1/2} y) \quad \text{otherwise}$$

we obtain the following expression:

$$u(x,y) = \sum_k u_k \, \varphi_k(x) \psi_k(y)$$

where the coefficients u_k are to be determined through (9c). Since we are dealing with smooth functions, the left-hand side of (9c) reduces to the limit $\lim_{y \downarrow 0} -y^\alpha u_y$. In the case of $s = 1/2$, we immediately obtain that:

$$u_y(x,0) = \sum_k u_k \, \varphi_k(x) \psi_k'(0) = -\sum_k u_k \mu_k^{1/2} \, \varphi_k(x)$$

Since $\{\varphi_k\}$ is a complete orthogonal set for $L^2(\Omega)$, we also have that $f = \sum_k f_k \varphi_k$. Equation (9c) implies that $f_k = u_k \mu_k^{1/2}$ for all k's and, thus, $(-\Delta)^{1/2} u = f$ on Ω as desired. In the case $s \neq 1/2$, in order to make use of equation (9c), we employ the following properties of function K_s:

1. $K_s(z) = K_{-s}(z)$;
2. $K_s'(z) = -\frac{1}{2}\big[K_{s-1}(z) + K_{s+1}(z)\big]$;
3. $K_s(z) \simeq \frac{1}{2}\Gamma(s)\big(\frac{1}{2}z\big)^{-s} \quad , \quad z \downarrow 0 \quad , \quad s > 0.$

We have that:

$$\begin{aligned}
\lim_{y\downarrow 0} y^\alpha \psi_k' &= \lim_{y\downarrow 0} y^\alpha \big[s(\mu_k^{1/2} y)^{s-1} \mu_k^{1/2} K_s(\mu_k^{1/2} y) + (\mu_k^{1/2} y)^s \mu_k^{1/2} K_s'(\mu_k^{1/2} y) \big] \\
&= \lim_{y\downarrow 0} y^\alpha \Big[s \mu_k^{s/2} y^{s-1} K_s(\mu_k^{1/2} y) - \mu_k^{(s+1)/2} y^s \tfrac{1}{2} \big[K_{s-1}(\mu_k^{1/2} y) + K_{s+1}(\mu_k^{1/2} y) \big] \Big] \\
&= \lim_{y\downarrow 0} y^\alpha \Big[s \mu_k^{s/2} y^{s-1} K_s(\mu_k^{1/2} y) - \mu_k^{(s+1)/2} y^s \tfrac{1}{2} \big[K_{1-s}(\mu_k^{1/2} y) + K_{s+1}(\mu_k^{1/2} y) \big] \Big] \\
&= \lim_{y\downarrow 0} y^\alpha \Big[s \mu_k^{s/2} y^{s-1} \tfrac{1}{2} \Gamma(s) \big(\tfrac{1}{2}\mu_k^{1/2} y\big)^{-s} + \\
&\qquad - \mu_k^{(s+1)/2} y^s \tfrac{1}{2} \big[\tfrac{1}{2} \Gamma(1-s) \big(\tfrac{1}{2}\mu_k^{1/2} y\big)^{s-1} + \tfrac{1}{2} \Gamma(s+1) \big(\tfrac{1}{2}\mu_k^{1/2} y\big)^{-s-1} \big] \Big] \\
&= \lim_{y\downarrow 0} y^\alpha \Big[\Gamma(s+1) \tfrac{1}{2}^{1-s} y^{-1} + \\
&\qquad - \Gamma(1-s) \tfrac{1}{2}^{s+1} \mu_k^s y^{2s-1} - \Gamma(s+1) \tfrac{1}{2}^{1-s} y^{-1} \Big] \\
&= \lim_{y\downarrow 0} y^{1-2s} \big[- \Gamma(1-s) \tfrac{1}{2}^{s+1} \mu_k^s y^{2s-1} \big] \\
&= -\Gamma(1-s) \tfrac{1}{2}^{s+1} \mu_k^s \\
&= -d_s \mu_k^s
\end{aligned}$$

where we defined $d_s := \Gamma(1-s) \frac{1}{2}^{s+1}$. Therefore equation (9c) can be equivalently written as

$$\lim_{y\downarrow 0} y^\alpha \sum_k u_k \varphi_k(x) \psi_k'(y) = d_s \sum_k u_k \mu_k^s \, \varphi_k(x) = \sum_k f_k \, \varphi_k(x)$$

where, by virtue of the fact that $\{\varphi_k\}$ is a complete set, we have that $f = \sum_k f_k \varphi_k$. Thus, $f_k = d_s u_k \mu_k^s$ for all k's and $(-\Delta)^s u = d_s f$ on Ω.

References

1. D.E. Amos. Algorithm 644: A portable package for Bessel functions of a complex argument and nonnegative order. *ACM Trans. Math. Software*, 12(3):265273, 1986. Single and double precision, maximum accuracy 18S.
2. P.W. Bates. On some nonlocal evolution equations arising in materials science. *Nonlinear dynamics and evolution equations*, 48:13–52, 2006.
3. J. Blackledge and M. Blackledge. Fractional anisotropic diffusion for noise reduction in magnetic resonance images. *ISAST Transactions on Electronics and Signal Processing*, 4:44–57, 2010.
4. C. Brändle, E. Colorado, and A. de Pablo. A concave-convex elliptic problem involving the fractional Laplacian. *ArXiv e-prints*, June 2010.
5. L. Caffarelli and L. Silvestre. An extension problem related to the fractional laplacian. *Communications in Partial Differential Equations*, 32(8):1245–1260, August 2007.
6. A. Capella, J. Dávila, L. Dupaigne, and Y. Sire. Regularity of radial extremal solutions for some non local semilinear equations. *ArXiv e-prints*, April 2010.
7. P. Carr, H. Geman, D.B. Madan, and M. Yor. The fine structure of asset returns: An empirical investigation. *Journal of Business*, 75:302–332, 2002.
8. L. Demkowicz. *Computing with hp-adaptive finite elements. Vol. 1*, volume 12 of *Chapman & Hall/CRC Applied Mathematics and Nonlinear Science Series*. Chapman & Hall/CRC, Boca Raton, FL, 2007.
9. L. Demkowicz and J. Kurtz. Projection-based interpolation and automatic *hp*-adaptivity for finite element discretizations of elliptic and maxwell problems. 21(September):1–15, 2007.
10. L. Demkowicz, W. Rachowicz, and Ph. Devloo. A fully automatic hp-adaptivity. *J. Sci. Comput.*, 17(1-4):117–142, December 2002.
11. A. Gelb and D. Gottlieb. The resolution of the Gibbs phenomenon for "spliced" functions in one and two dimensions. *Computers & Mathematics with Applications*, 33(11):35–58, June 1997.
12. V. Gol'dshtein and A. Ukhlov. Weighted Sobolev spaces and embedding theorems. *Transactions of the American Mathematical Society*, 361(7):3829–3850, 2009.
13. Q.Y. Guan and Z.M. Ma. Reflected symmetric α-stable processes and regional fractional laplacian. *Probability theory and related fields*, 134(4):649–694, 2006.
14. L. Hörmander. *The Analysis of Linear Partial Differential Operators III: Pseudo-Differential Operators*. A Series of Comprehensive Studies in Mathematics / Grundlehren Der Mathematischen Wissenschaften. Springer, 1985.
15. J.L. Lions and E. Magenes. *Non-homogeneous boundary value problems and applications*. Non-homogeneous Boundary Value Problems and Applications. Springer-Verlag, 1972.
16. R.L. Magin, C. Ingo, L. Colon-Perez, W. Triplett, and T.H. Mareci. Characterization of anomalous diffusion in porous biological tissues using fractional order derivatives and entropy. *Microporous and Mesoporous Materials*, 178:39–43, 2013.
17. W.C.H. McLean. *Strongly Elliptic Systems and Boundary Integral Equations*. Cambridge University Press, 2000.
18. R.H. Nochetto, E. Otarola, and A.J. Salgado. A PDE approach to fractional diffusion in general domains: a priori error analysis. *arXiv.org*, February 2013.
19. L. Silvestre. *Regularity of the obstacle problem for a fractional power of the Laplace operator*. PhD thesis, May 2005.
20. E. Stein. *Singular integrals and differentiability properties of functions*. Princeton Mathematical Series, No. 30. Princeton University Press, Princeton, N.J., 1970.
21. P.R. Stinga and J.L. Torrea. Extension problem and Harnack's inequality for some fractional operators. *ArXiv e-prints*, October 2009.
22. J.L. Vázquez. Recent progress in the theory of Nonlinear Diffusion with Fractional Laplacian Operators. *arXiv preprint arXiv:1401.3640*, pages 1–33, January 2014.

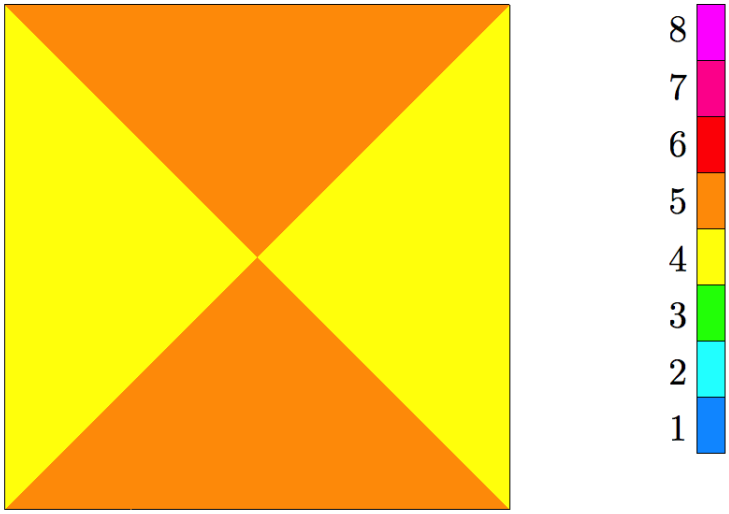

(a) Optimal mesh ($\varepsilon = 0.0088$), $\#\,\mathrm{dof} = 30$.

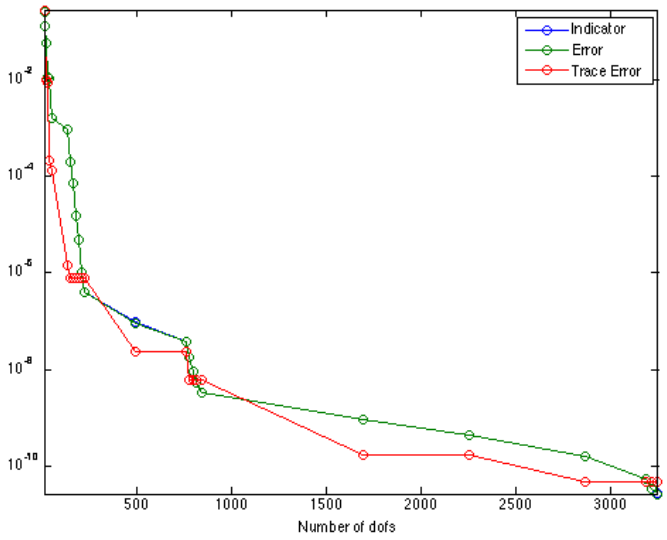

(b) Convergence rates (the Error curve overlays the Indicator curve).

Fig. 1 Case $s = 1/2$. The colors of the optimal mesh represent the orders of approximation of the element, according to the color scale on the right. The optimal mesh is comprised a single element with horizontal order of approximation equal to five, and vertical order of approximation equal to four.

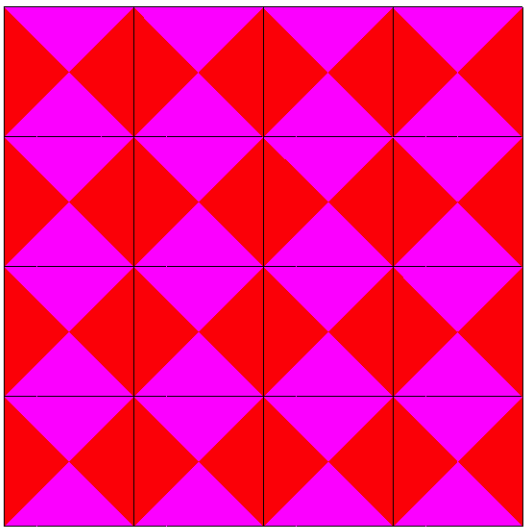

(a) Optimal mesh ($\varepsilon = 0.008$), $\#\,\mathrm{dof} = 165$.

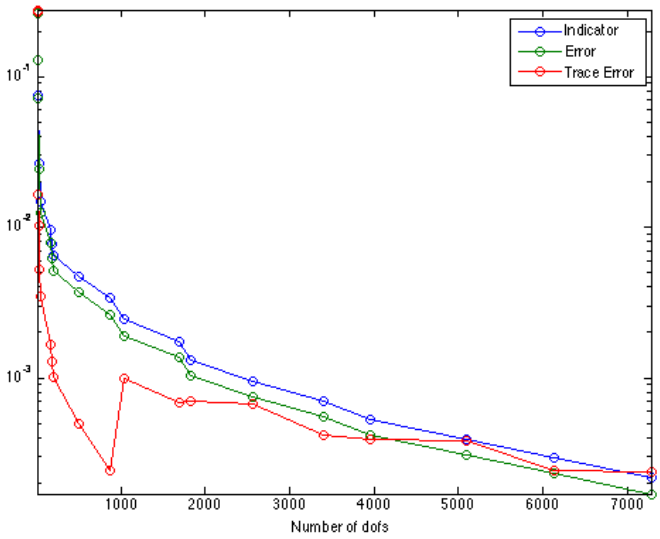

(b) Convergence rates.

Fig. 2 case $s = 0.48$. The colors of the optimal mesh represent the orders of approximation of the elements, see Figure 1.

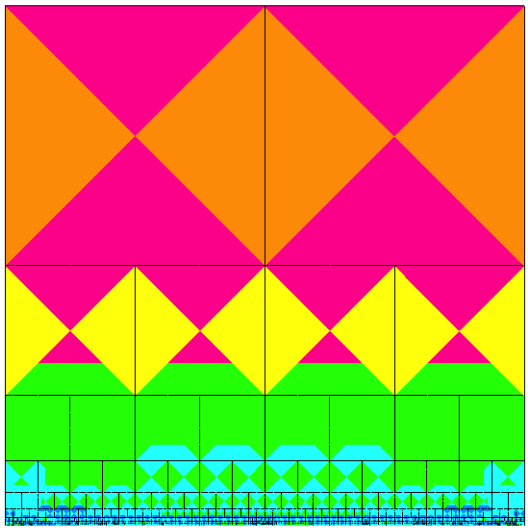

(a) Optimal mesh ($\varepsilon = 0.0091$), $\#\,\mathrm{dof} = 5817$.

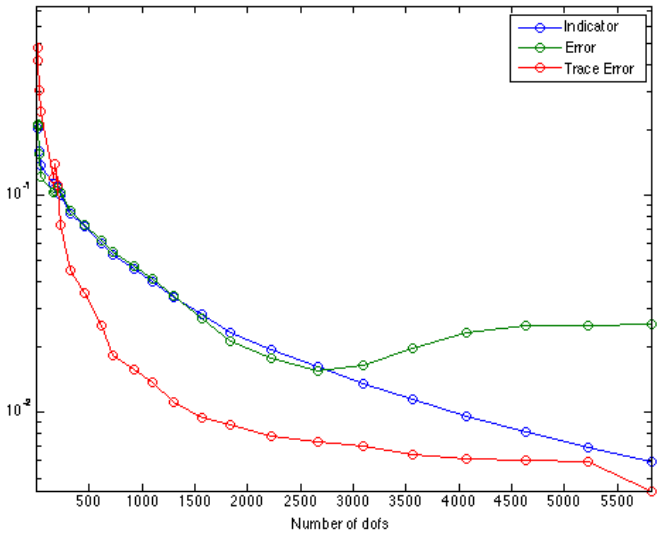

(b) Convergence rates.

Fig. 3 Case $s = 0.25$. The colors of the optimal mesh represent the orders of approximation of the elements, see Figure 1.

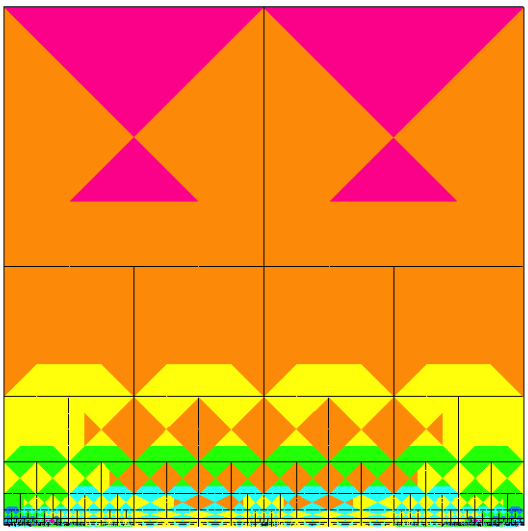

(a) Optimal mesh ($\varepsilon = 0.1089$), $\#\,\mathrm{dof} = 13576$.

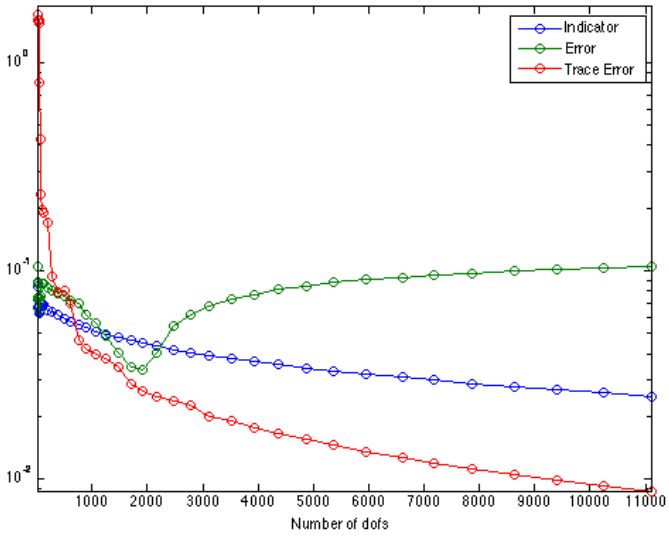

(b) Convergence rates.

Fig. 4 Case $s = 0.05$. The colors of the optimal mesh represent the orders of approximation of the elements, see Figure 1.

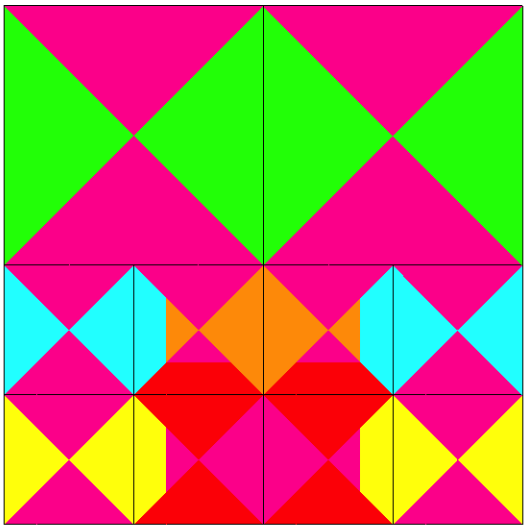

(a) Optimal mesh ($\varepsilon = 0.0097$), $\#\,\mathrm{dof} = 296$.

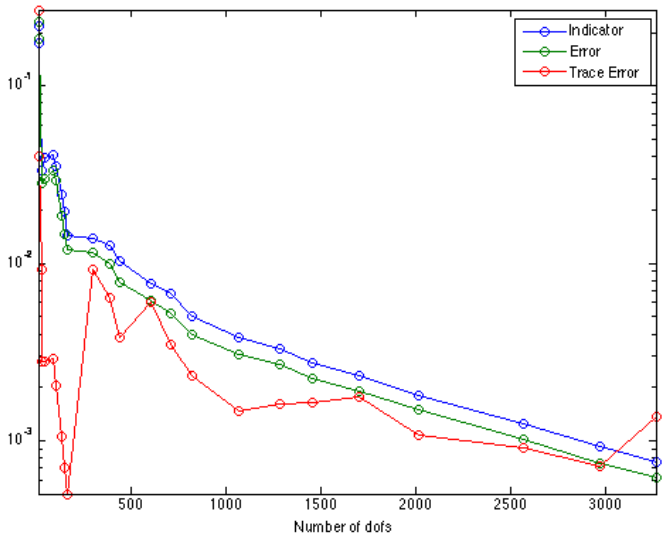

(b) Convergence rates.

Fig. 5 Case $s = 0.55$. The colors of the optimal mesh represent the orders of approximation of the elements, see Figure 1.

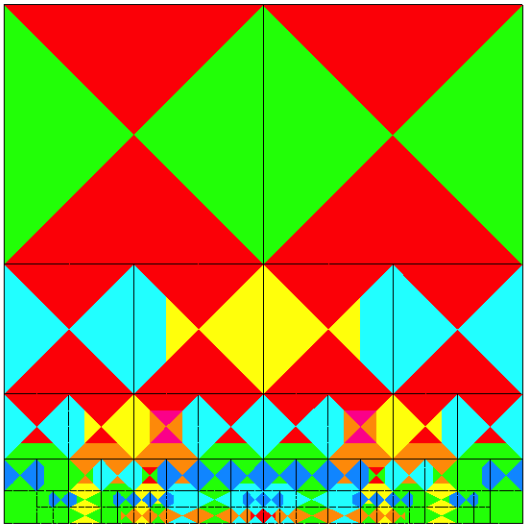

(a) Optimal mesh ($\varepsilon = 0.0083$), $\#\,\mathrm{dof} = 774$.

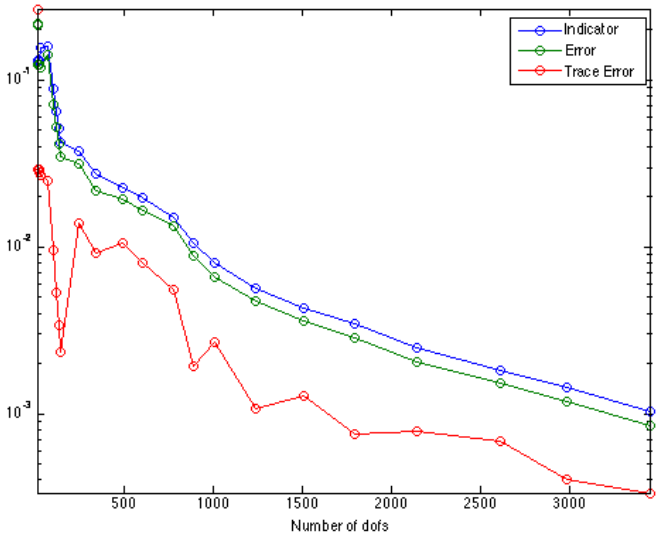

(b) Convergence rates.

Fig. 6 Case $s = 0.65$. The colors of the optimal mesh represent the orders of approximation of the elements, see Figure 1.

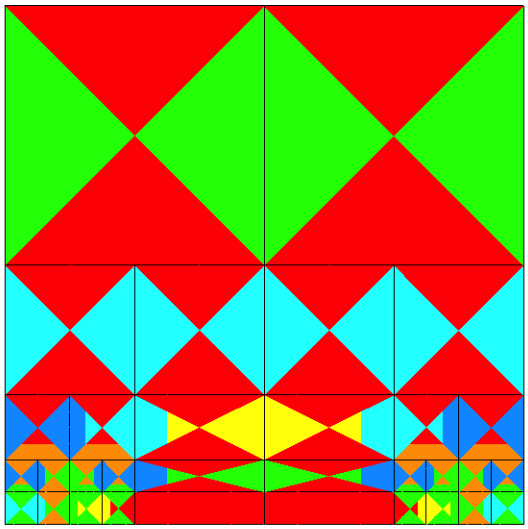

(a) Optimal mesh ($\varepsilon = 0.0080$), $\#\,\mathrm{dof} = 451$.

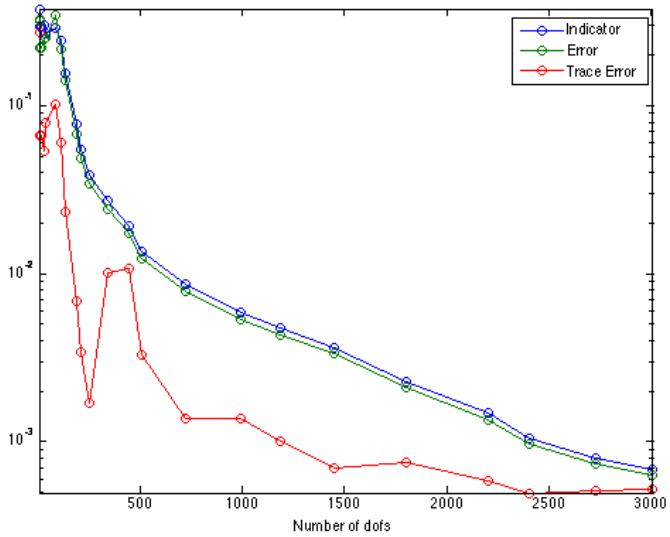

(b) Convergence rates.

Fig. 7 Case $s = 0.75$. The colors of the optimal mesh represent the orders of approximation of the elements, see Figure 1.

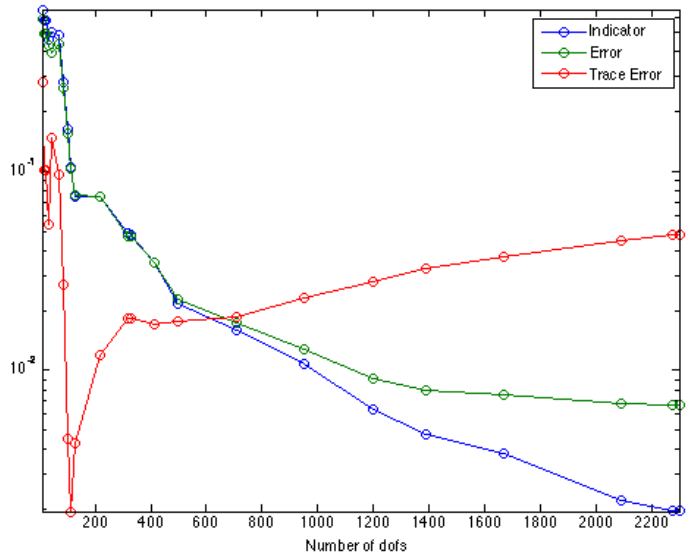

Fig. 8 Case $s = 0.85$. Convergence rates.

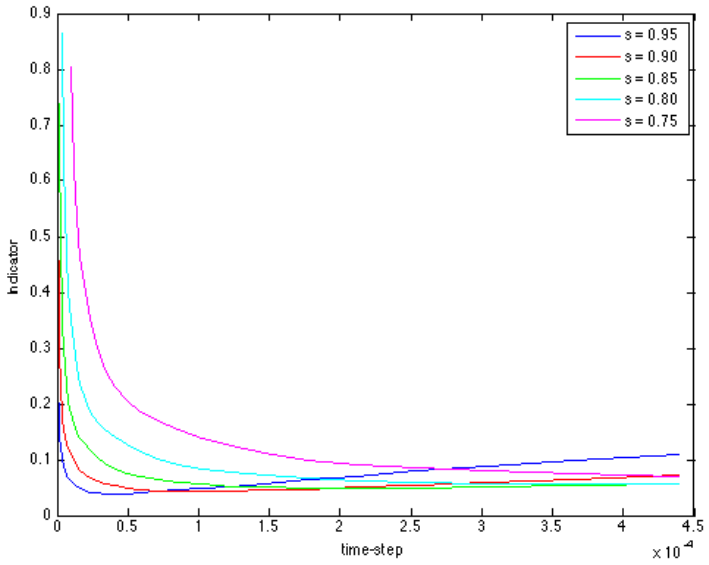

Fig. 9 Behavior of indicator E defined by (8) for different values of s. As s decreases, the tail of E becomes increasingly flat.

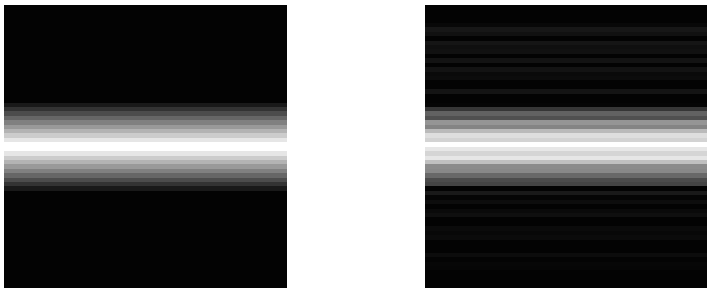

(a) Original image. (b) Corrupted image.

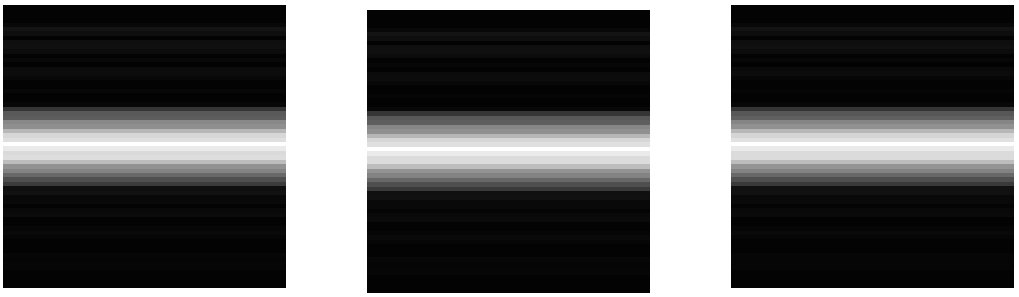

(c) Step 2: rel. error = 48.55% (d) Step 4: rel. error = 48.53% (e) Step 6: rel. error = 48.49%

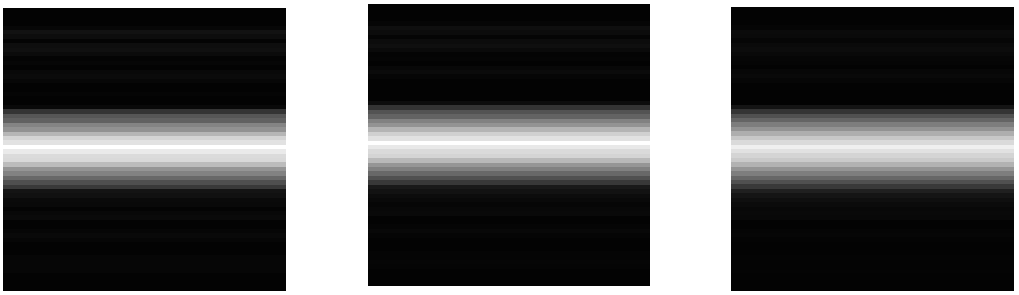

(f) Step 8: rel. error = 48.39% (g) Step 10: rel. error = 47.88% (h) Step 12: rel. error = 46.58%

Fig. 10 Adaptive denoising strategy of Algorithm 1 applied to a unit spike function. The signal was corrupted by a random noise uniformly distributed on $[-0.1, 0.1]$. Twelve steps were required for the denoising process. Gray-scale images, whose vertical cross-sections coincide with the original, noisy and denoised functions are shown. While an improvement at each step is visible with respect to the "eyeball" norm, the decrease of the L^1 relative error is only minimal.

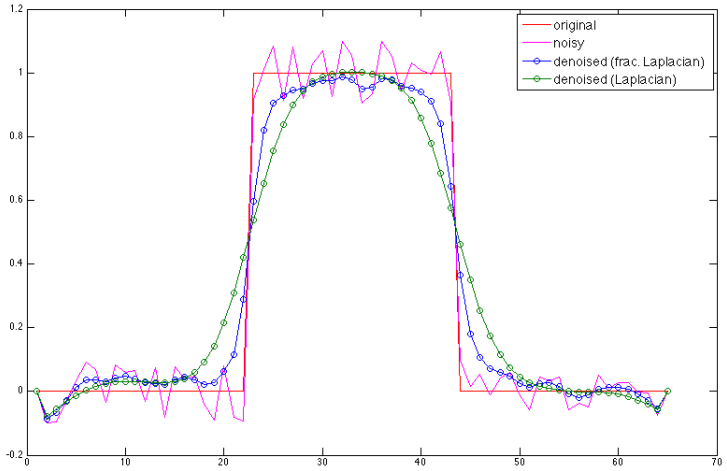

(a) Adaptive denoising of a unit step function.

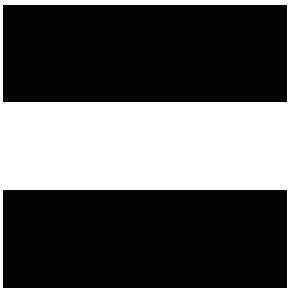

(b) Original image

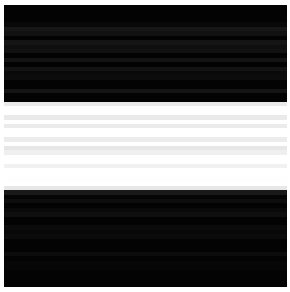

(c) Corrupted image

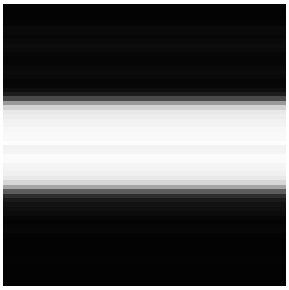

(d) Denoised image (fractional Laplacian).

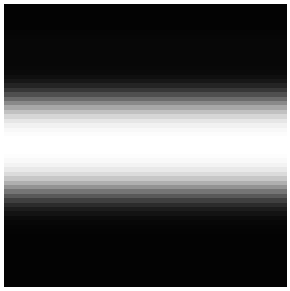

(e) Denoised image (standard Laplacian).

Fig. 11 Adaptive denoising strategy of Algorithm 1 applied to a unit step function. The signal was corrupted by a random noise uniformly distributed on $[-0.1, 0.1]$. Gray-scale images, whose vertical cross-sections coincide with the original, noisy and denoised functions are shown in order to assess the quality of the result in the "eyeball" norm.

(a) original image

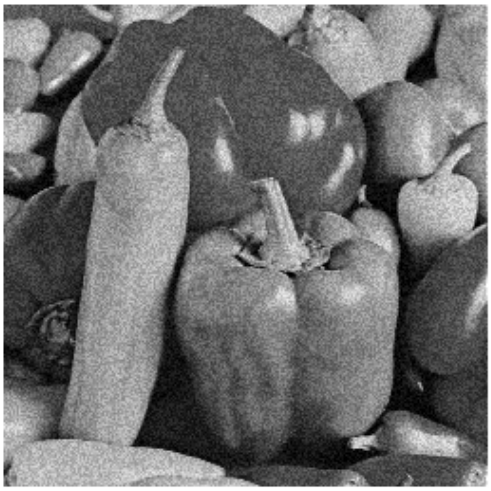

(b) noisy image

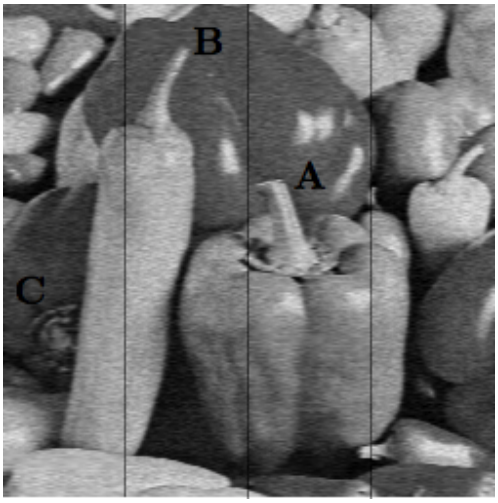

(c) denoised image (Laplacian)

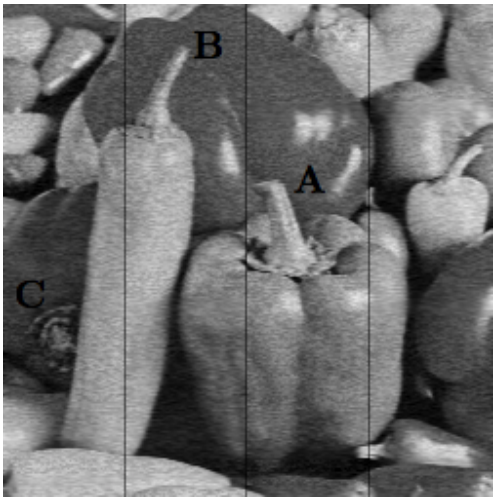

(d) denoised image (fractional Laplacian)

Fig. 12 "Peppers" image denoising. The original image (top left) was corrupted by adding a uniformly distributed random noise (top right). The image was cut vertically into four parts, and each sub-image was processed independently by applying the standard Laplacian (bottom left) and fractional Laplacian (bottom right). Pairwise comparison of details from the processed images shows that the fractional Laplacian preserves sharper details than the standard Laplacian. No Gibbs phenomenon can be observed on the vertical cuts.

(a) original image

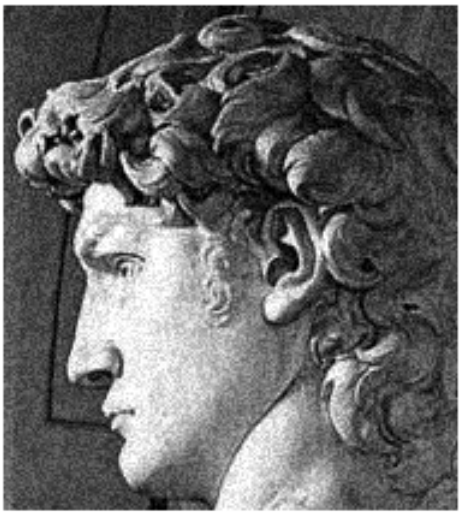

(b) noisy image

(c) denoised image (Laplacian)

(d) denoised image (fractional Laplacian)

Fig. 13 David image denoising. The original image (top left) was corrupted by adding a uniformly distributed random noise (top right). The image was cut vertically into three parts, and each sub-image was processed independently by applying the standard Laplacian (bottom left) and fractional Laplacian (bottom right). Comparison of the processed images shows that the fractional Laplacian preserves sharper details than the standard Laplacian. No Gibbs phenomenon can be observed on the vertical cuts.