AN RBF PARTITION OF UNITY METHOD FOR GEOMETRY RECONSTRUCTION AND PDE SOLUTION IN THIN STRUCTURES

ELISABETH LARSSON^{*}, PIERRE-FRÉDÉRIC VILLARD[†], IGOR TOMINEC[‡], ULRIKA SUNDIN^{*}, ANDREAS MICHAEL^{*}, AND NICOLA CACCIANI[§]

Abstract. The main respiratory muscle, the diaphragm, is an example of a thin structure. We aim to perform detailed numerical simulations of the muscle mechanics based on individual patient data. This requires a representation of the diaphragm geometry extracted from medical image data. We design an adaptive reconstruction method based on a least-squares radial basis function partition of unity method. The method is adapted to thin structures by subdividing the structure rather than the surrounding space, and by introducing an anisotropic scaling of local subproblems. The resulting representation is an infinitely smooth level set function, which is stabilized such that there are no spurious zero level sets. We show reconstruction results for 2D cross sections of the diaphragm geometry as well as for the full 3D geometry. We also show solutions to basic PDE test problems in the reconstructed geometries.

Key words. radial basis function, least-squares, partition of unity, implicit surface, RBF-PUM

AMS subject classifications. 65N35, 65N12

1. Introduction. This work has been performed with a particular application in mind. We aim to build a realistic model of the human respiratory system, with a special focus on the respiratory muscles, to enable in silico investigation of the adverse effects of mechanical ventilation on the muscle tissue, as well as the benefits of potential improvements to the mechanical ventilation process [17]. To meet the specific challenges of this application we develop numerical tools that also have a wider use. The types of problems that we consider have the following characteristics:

- The problem can be formulated as (a system of) partial differential equations (PDE). Furthermore, PDE coefficients, boundary data, and, if present, initial values, are smooth and compatible such that they admit a smooth solution.
- The PDE is formulated in a smooth, not easily parametrized geometry.
- The geometry has a high aspect ratio. For the diaphragm, which is the main respiratory muscle, the ratio is approximately 1:100.

The simulation probem consists of two main parts. First, the geometry needs to be represented in such a way that its properties, such as boundary location and surface normal, can be incorporated into the PDE formulation. Then, the PDE is solved, to generate data of interest for the application specialist. The assumed smoothness, present also for other biomechanical problems, suggests that we may benefit from using a high order method. Coupled with a non-trivial geometry, methods based on radial basis functions (RBF) have an advantage regarding the flexibility and ease of implementation. In this paper, we have chosen to use an unfitted least-squares RBF partition of unity method [23] both for the reconstruction and for the PDE solution to create a unified framework for the whole simulation problem.

^{*}Uppsala University, Department of Information Technology, SE-751 05 Uppsala, Sweden (elisabeth.larsson@it.uu.se, ulrika.sundin@it.uu.se, andreas.michael@it.uu.se).

[†]Université de Lorraine, CNRS, Inria, LORIA, Nancy, France (pierrefrederic.villard@loria.fr).

 $^{^{\}ddagger}$ Stockholm University, Department of Mathematics, SE-106 91 Stockholm, Sweden (igor.tominec@math.su.se).

[§]Basic and Clinical Muscle Biology, Department of Physiology and Pharmacology and Center for Molecular Medicine, Karolinska Institutet, Stockholm, Sweden (nicola.cacciani@ki.se).

Funding: This work was supported by the Swedish Research Council, grants no. 2016-04849 and no. 2020-03488.

In our previous work on the diaphragm application [2, 36, 38], we used the RBF-FD method [12] and an unfitted least-squares RBF-FD method [37, 36], that has improved numerical stability properties in the presence of derivative boundary conditions compared with collocation RBF-FD. In [2] and [38], linear elasticity model problems were solved and we investigated the need for resolution of the thin dimension and how to handle mixed boundary conditions in a smooth setting, respectively. In [36] the focus is on the unfitted method, Possion model problems are solved, and the diaphragm geometry is used as a demonstrator.

Reconstruction of a thin structure has many similarities with surface reconstruction, but with the added complication of interference between the two nearby surfaces [1]. RBF-based methods (kernel methods) have been used for surface reconstruction for more than two decades [3, 29]. As in these references, we focus on implicit surface representations, where a level set function determines which points are on the surface (zero level set), outside (positive values), or inside (negative values). Due to the typically large data sets, localized methods with low computational complexity are preferred. Popular approaches include using compactly supported RBFs [28], using a partial of unity method combined with an adaptive oct tree subdivision of space [27, 44, 45], or using quasi interpolation [25]. Least-squares approximation, regularization, and smoothing are used to suppress noise in the data [4, 27, 45, 25]. Data on the surface is often complemented by exterior and interior data points to guide the direction of the level set function [3, 29], whereas our work is inspired by [30], where gradient information is incorporated using a Hermite type RBF approximation. In the recent paper [9], the gradient field is instead approximated by a curl-free RBF partition of unity method, which then implicitly provides the level set function. Another recent paper [42], considers reconstruction of thin plant leaves. However, in this case the aspect ratio is high enough that the leaves are considered as surfaces.

Our objective is to create an infinitely (or high-order) smooth geometry reconstruction in the form of an implicit surface representation, starting from a noisy point cloud representation. Furthermore, to facilitate the use of the geometry in the PDE problem, we require the level set function to be free from spurious zero level sets inside and outside the geometry. We do this using the least-squares RBF partition of unity method (RBF-PUM) from [23], where node generation is decoupled from the geometry of the computational domain, with the following specific adaptions:

- The thin structure is adaptivily subdivided instead of the space.
- The same type of gradient conditions as in [30], but without the Hermite basis, are combined with exterior and interior stabilizing conditions.
- The risk of interference between nearby surfaces is reduced by anisotropic scaling of the local subproblems. For a somewhat related approach, see [5].

Least-squares RBF-PUM was applied to elliptic PDEs in [23]. Model problems with parametrized geometries were solved. Furthermore, a uniform grid-based distribution of spherical patches covering the domain was used. In this paper, we evaluate the performance of least-squares RBF-PUM combined with our smooth geometry reconstruction and with a non-uniform distribution of cylindrical patches. To achieve convergence in a finite precision setting when using infinitely smooth RBFs, a stable evaluation method, such as the RBF-QR method [14, 13, 20], is needed. The 3D RBF-QR implementation [20] previously did not contain all second derivative operators. For the experiments in this paper, we have updated the method with the full set [22]. As for the reconstruction problem, in the PDE solver, we allow anisotropic scaling of the local subproblems. However, in this case the scaling is employed to balance the resolution across the thin dimension against the resolution along the other two dimensions.

A convergence estimate for least-squares RBF-PUM was derived in [23]. The estimate contains a stability norm, that was investigated numerically, and the interpolation error, for which theoretical estimates are available. The convergence rate is provided by the interpolation error and we could show numerically that the stability norm for least-squares RBF-PUM stays constant under patch refinement, which is not the case for collocation RBF-PUM. Recently, in [37], we used a different technique to analyze the stability of a least-squares RBF-FD method. This allowed us to prove stability, and we plan to revisit the analysis of least-squares RBF-PUM using this technique. In this paper, we adapt the interpolation error estimates to our specific setting, to explain the observed convergence rates.

The outline of the paper is as follows. In section 2 we describe the data sets we use for the geometry reconstruction. Section 3 briefly introduces the least-squares RBF-PUM method. Section 4 and section 5 provide details on how we generate the partition of unity cover and weight functions. Section 6 describes how we use leastsquares RBF-PUM for geometry reconstruction with results in section 7. Section 8 describes how the method is applied to the PDE problem with numerical results in section 9. The paper ends with conclusions in section 10.

2. The noisy input geometry data. For our target application, the input data is extracted from 3D medical images. The extraction process is partly based on manual segmentation using physiological knowledge [39], and partly automatic processing using the Marching Cube algorithm to create a mesh representing the geometry, followed by the decimation method in [24] to generate the final data set. These steps are described in more detail in [2]. The data consists of a point cloud Y = $\{y_k\}_{k=1}^M \in \mathbb{R}^d$, for d=2 or d=3, representing the curve/surface and approximate surface normals \underline{n}_k , $k = 1, \ldots, M$ at the corresponding points. These algorithms generate data that is quasi uniformly distributed over the surface. For simplicity, we assume this property in the algorithms that are derived in the following sections, but non-uniform distributions could also be taken into account.

In addition to the data, we make use of the initial approximate surface representation as a mesh provided by the extraction algorithm. If an initial mesh representation was not available, we could instead create local triangulations, on demand, that would serve our purposes.

An example of a point cloud representing a diaphragm geometry is shown in the left part of Figure 2.1. In the numerical experiments we also use a 2D cross section of the diaphragm. This data is shown in the right part of the figure. It is visually clear from the normals in the 2D case that the data is noisy and contains artifacts from the manual segmentation, taking the form of bumps.

3. Least-squares RBF-PUM. In a partition of unity (PU) method, the computational domain $\Omega \subset \mathbb{R}^d$ is covered by a set of overlapping patches $\{\Omega_j\}_{j=1}^P$, such that $\bar{\Omega} \subset \bigcup_{j=1}^{P} \bar{\Omega}_j$. We require each point $\underline{x} = (x_1, \ldots, x_d) \in \bar{\Omega}$ to be in the interior of at least one patch. We also require the overlap to be mild in the sense that no more than K patches overlap at any given point. We construct a set of non-negative partition of unity weight functions $\{w_j(\underline{x})\}_{j=1}^{P}$, where w_j is compactly supported on Ω_j . The partition of unity property implies that $\sum_{j=1}^P w_j(\underline{x}) = 1$, for $\underline{x} \in \bigcup_{j=1}^P \Omega_j$. A global PU approximation $\tilde{u}(\underline{x})$ is constructed from local approximations $\tilde{u}_j(\underline{x})$



FIG. 2.1. The initial point cloud data in 3D (left) and 2D (right), together with the initial normal data. The surface is divided into an inner and outer part whose normals are displayed in different colors.

on each patch and the weight functions as

(3.1)
$$\tilde{u}(\underline{x}) = \sum_{j=1}^{P} w_j(\underline{x}) \tilde{u}_j(\underline{x}).$$

The smoothness of the global approximation is determined by the smoothness of the weight functions and of the local approximations. In this paper, we use RBF-PUM both for the geometry representation and for solving a PDE problem formulated in that geometry. For the PDE problem it is enough if the smoothness of the weight function supports the differential operators of the PDE. However, to construct a high-order method for the PDE problem, we need high-order smoothness of the geometry. That is, for the geometry representation, we require high-order smoothness also of the weight functions.

For the local RBF approximations, we make use of a template node layout in a reference patch Ω_0 , see [23]. The patches Ω_j are allowed to have different sizes, different aspect ratios, and different orientations. For each patch, we introduce a map $T_j : \Omega_j \mapsto \Omega_0$ that takes points in Ω_j into the reference patch. We let the local approximations be given by

(3.2)
$$\tilde{u}_j(\underline{x}) = \tilde{v}_j(T_j(\underline{x})) \triangleq \tilde{v}_j(\underline{x}'),$$

where $\tilde{v}_j(\underline{x}')$ is an RBF approximation in the reference patch, of the form

(3.3)
$$\tilde{v}_j(\underline{x}') = \sum_{i=1}^n \lambda_i^j \phi(\|\underline{x}' - \underline{x}_i'\|) \triangleq \sum_{i=1}^n \lambda_i^j \phi_i(\underline{x}'),$$

where λ_i^j are unknown coefficients, $\phi(r)$ is an RBF, and $X' = \{\underline{x}_i'\}_{i=1}^n$ is the set of RBF center points in the reference patch, with fill distance h'. The corresponding points in Ω_j are $\underline{x}_i^j = T_j^{-1}(\underline{x}_i')$, $i = 1, \ldots, n$, with fill distance h. We let $\underline{u}_j = (\tilde{u}_j(\underline{x}_1^j), \ldots, \tilde{u}_j(\underline{x}_n^j))^{\mathsf{T}} = (\tilde{v}_j(\underline{x}_1'), \ldots, \tilde{v}_j(\underline{x}_n'))^{\mathsf{T}}$ and $\underline{\lambda}_j = (\lambda_1^j, \ldots, \lambda_n^j)^{\mathsf{T}}$. Then, using (3.3), we have $\underline{u}_j = A\underline{\lambda}_j$, where the local interpolation matrix A has elements $a_{ki} = \phi(\|\underline{x}_k' - \underline{x}_i'\|)$. For commonly used infinitely smooth RBFs such as the multiquadric with $\phi(r) = \sqrt{1+r^2}$, the Gaussian with $\phi(r) = e^{-r^2}$, and the inverse multiquadric with $\phi(r) = 1/\sqrt{1+r^2}$, the interpolation matrix A is non-singular [34, 26]. This implies that we can safely use $\underline{\lambda}_j = A^{-1}\underline{u}_j$ in (3.3) to get the local RBF approximations expressed in terms of the nodal values as

(3.4)
$$\tilde{v}_j(\underline{x}') = (\phi_1(\underline{x}'), \dots, \phi_n(\underline{x}'))A^{-1}\underline{u}_j.$$

Note that in this formulation of RBF-PUM, each local approximation has its own nodal values, also in the overlapping regions. The total number of nodal unknowns to determine is N = nP.

When we use RBF-PUM to approximate a function, the local subproblems can be solved independently, and are then blended into the global approximation. For a PDE, we instead need to assemble a global system of equations to get a well-posed discretized problem. We describe a generic time-independent linear PDE by

(3.5)
$$\mathcal{L}(\underline{x})u(\underline{x}) = f(\underline{x}), \quad \underline{x} \in \overline{\Omega}.$$

where $\mathcal{L}(x)$ represents the differential operators of the PDE and its boundary conditions, and f(x) represents all right hand side functions. To form the global leastsquares system of equations, we sample the PDE problem at the points $\{\underline{y}_k\}_{k=1}^M$, where M > N. If we let $\mathcal{L}_k = \mathcal{L}(\underline{y}_k)$, then the equations are given by

(3.6)
$$\mathcal{L}_k \tilde{u}(\underline{y}_k) = \sum_{j=1}^P \mathcal{L}_k(w_j(\underline{y}_k)\tilde{u}_j(\underline{y}_k)) = f(\underline{y}_k)$$

Practically, the matrix is assembled by adding the contributions from each patch. We write the global system as $L\underline{u} = \underline{f}$, where the vector $\underline{u} = (\underline{u}_1^{\mathsf{T}}, \ldots, \underline{u}_P^{\mathsf{T}})^{\mathsf{T}}$ and the vector $\underline{f} = (f(\underline{y}_1), \ldots, f(\underline{y}_M))^{\mathsf{T}}$. Note that we differentiate the product of the weight functions and the local approximations such that each differential operator is expanded into several terms.

When solving the least squares system of equations, the scaling of the equations corresponding to different operators becomes important for convergence. Here, we use the insights from [37], where we dereived and analyzed a least-squares RBF-FD method. First we group the equations according to the type of operator. If there are M_i equations of type i, then we scale these equations with $1/\sqrt{M_i}$. With this scaling, the discrete least-squares problem is a discrete quadrature approximation of the corresponding continuous least-squares problem. Then, we increase the weight of the lower order boundary equations by multiplying these with a small negative power of the fill distance h of the reference node set X. This improves stability without affecting the rate of convergence.

4. Adaptive cover generation for thin structures. Common choices of patch shapes for RBF-PUM in the literature are discs and spheres or their anisotropic counterparts, see, e.g., [32, 23, 8, 7, 6]. These work well when the computational domain is isotropic or has a fixed direction of anisotropy. Here, we are considering thin curved structures, and it becomes relevant to use a patch shape that adheres to this property. We have chosen to use cylindrical patches that are radial in directions tangential to the surface and with a height that is related to the thickness of the domain. We also make the assumption that there is only one layer of patches in the thickness direction. We let the reference patch, centered at the origin, have radius R_0 and height H_0 . An example of two-dimensional cylindrical patches can be seen in Figure 4.1 and three-dimensional cylindrical patches are shown in Figure 4.2.

The starting point for the adaptive cover generation is the input data Y with M points and the user defined parameter P_0 corresponding to the initial number of patches. After the adaptive refinement, the final number of patches is $P \ge P_0$. The cover algorithm consists of the following steps:

1. Generate and adaptively refine patches until a volume criterion is met.

- (a) Use k-means clustering to form non-overlapping subsets $Y_j = \{\underline{y}_i^j\}_{i=1}^{m_j}$ of the input data with corresponding patch centers \underline{C}_j .
- (b) Use principal component analysis (PCA) on each Y_j to find the patch radius R_j , the vertical range $[Z_{j,1} Z_{j,2}]$, and the transformation T_j to the reference patch.
- (c) Evaluate the volume criterion, flag patches for refinement that fail the criteria, and return to 1(a) if not done.
- 2. Ensure that each boundary segment or surface triangle is entirely covered.
- 3. Ensure an overlap δ between the patches.

The volume criterion that we use in Step 1 aims to have a similar resolution (fill distance) in each patch. This also means that we refine where the curvature of the domain is large within a patch, since this increases the volume. We denote the initial mean volume by \bar{V}_0 and the initial sample standard deviation by s_{V_0} , and set the (absolute) target volume V^* as

(4.1)
$$V^* = \bar{V}_0 + \max(\bar{V}_0/10, s_{V_0}),$$

where the constant 10 is chosen to prevent excessive refinement when patches are already similar in volume. When patches for which $V_i > V^*$ have been flagged for refinement, we first add also the neighbours of each of these patches to the flagged set. Then for each connected set of p flagged patches, we take the union of the corresponding p node subsets, change the required number of patches to p + 1 for this subset of the data and construct these p + 1 new patches locally. This process is illustrated by the first tree panels of Figure 4.1.



FIG. 4.1. The adaptive cover generation with the initial patches (top left), the intermediate result (top right), and the final non-overlapping patches (bottom left). In each step, patches that that are marked for refinement (dashed lines) and the previous iteration (dotted lines) are indicated. The data points (dot markers) are shown for reference. Finally, the resulting patches after ensuring an overlap $\delta \approx 0.15$ are shown (bottom right).

The first step of the patch construction for a node set is k-means clustering. The algorithm aims to create clusters that are of almost equal size and with a minimal sum of square distances to the centroid. An example of a clustering is shown in the left part of Figure 4.2. The initial patch centers \underline{C}_j are taken as the centroids of the clusters.

In the local patch coordinate system, we let \underline{C}_j be the origin. We let $\tilde{Y}_j = \{\underline{y}_i^j - \underline{C}_j\}_{i=1}^{m_j}$, and we denote the matrix of size $m_j \times d$ containing the points by



FIG. 4.2. The result of k-means clustering of 10276 nodes into 75 clusters (left) and the P = 79 final patches generated by the 3 steps of the cover algorithm (right). Only the mantle of the cylindrical patches is drawn to make the intersections visible.

 $[\tilde{Y}_j]$. A principal component analysis of the data amounts to finding the eigenvalues $\mu_1 \geq \cdots \geq \mu_d$ of $[\tilde{Y}_j]^{\intercal}[\tilde{Y}_j]$ and the corresponding eigenvectors q_1, \ldots, q_d . The first d-1 principal components (eigenvectors) represent the main components of the data, i.e., the tangent (plane) direction(s), while the *d*th eigenvector is the direction normal to the tangent plane.

We form the orthogonal matrix Q_j with the eigenvectors as columns. We transform the local points into the patch coordinate system and get $\hat{Y}_j = \{ \underline{\tilde{y}}_i^j Q_j \}_{i=1}^{m_j}$, which allows us to compute the radius and vertical range of the patch:

(4.2)
$$R_j = \max_i \left(\sum_{k=1}^{d-1} (\underline{\hat{y}}_{i,k}^j)^2 \right)^{\frac{1}{2}},$$

We denote the center of the patch in the local coordinate system by \underline{Z}_j , which has d-1 zero components and $\frac{1}{2}(Z_{j,1}+Z_{j,2})$ in the last position. We also define the diagonal scaling matrix S_j with diagonal elements $s_{ii} = \frac{R_0}{R_j}$ for $i = 1, \ldots, d-1$ and $s_{dd} = \frac{H_0}{H_j}$. The local points in the reference patch coordinate system then become $Y'_j = \{(\underline{\hat{y}}_i^j - \underline{\bar{Z}}_j)S_j\}_{i=1}^{m_j}$. We now have all the steps to write down the transformation from the global coordinate system to the reference patch for the points in patch j

(4.5)
$$T_j(\underline{x}) = ((\underline{x} - \underline{C}_j)Q_j - \underline{Z}_j)S_j = \underline{x}'.$$

When expressing conditions on functions in the physical domain in terms of the approximations in the reference patch, we need to differentiate this map. If we define the gradient as a row vector, then for a function $f(\underline{x}) = g(T_j(\underline{x})) = g(\underline{x}')$

(4.6)
$$\nabla f = \nabla' g S^{\mathsf{T}} Q^{\mathsf{T}},$$
$$\mathcal{H}_{f} = Q S \mathcal{H}'_{g} S^{\mathsf{T}} Q^{\mathsf{T}},$$
$$\Delta f = \operatorname{Tr}(S \mathcal{H}'_{g} S^{\mathsf{T}}) = \sum_{i=1}^{d} s_{ii}^{2} \frac{\partial^{2} g}{\partial x_{i}^{\prime 2}},$$

where \mathcal{H} denotes the Hessian matrix, primed operators refer to the reference coordinate system, and the Laplacian has been simplified using rotational invariance.

The patches that have been generated so far cover all data points, but to ensure a cover of the whole geometry, we also need to cover the potential gaps between the data points. Here we use the initial approximate surface representation, see section 2. We extend the patches such that each surface element is covered. The surface elements are planar objects, while patches are convex regions. This implies that if all corners of an element are covered by a patch, then the whole element is covered. For elements with their corners in different patches, we divide the element into d Voronoi regions associated with the d corners, and then extend the closest patch that covers a corner to cover also the Voronoi region of that corner. When all boundary elements are covered, we have ensured that the whole initial object is covered.

The final step of the cover algorithm ensures that there is sufficient overlap between the patches. The overlap is important because it affects the steepness of the weight functions. Across an overlap region, the involved weight functions need to transition smoothly from 0 to 1. To handle this practically, we require that all data points have at least a at relative distance δ_0 from the inside of the patch to any patch boundary. For each point where this is not satisfied, we extend the patch that requires the smallest volume increase for including that point. The resulting relative overlap will not be precisely δ_0 , but close enough as long as the data points are dense enough. The result of the overlap algorithm is shown in the last panel of Figure 4.1 for the two-dimensional diaphragm geometry.

Remark 4.1. If the surface data is too sparse in relation to the thickness of the geometry, the clustering can be performed only on the inner or outer surface in order to acheive the desired result with one layer of patches. We have implemented an automatic approach to label the edge, the inner surface, and the outer surface.

5. Tensor product partition of unity weight functions. The partition of unity weight functions are tightly connected with the cover. To build the weight functions, we start from a set of non-negative generating functions supported on the respective patches. The cylindrical patch shape that we have chosen suggests a tensor product form. Let $\psi_0(r)$ be a positive radial function supported on [0, 1), and let $\|\underline{x}\|_r = \sqrt{x_1^2 + \cdots + x_{d-1}^2}$. Then the product

(5.1)
$$\psi(\underline{x}') = \psi_0 \left(\frac{\|\underline{x}'\|_r}{R_0}\right) \psi_0 \left(\frac{2|x_d|}{H_0}\right)$$

is compactly supported on the reference patch with radius R_0 and height H_0 . The generating function for patch j is then given by

(5.2)
$$\psi_j(\underline{x}) = \psi(T_j(\underline{x})) = \psi(\underline{x}').$$

The generating functions we use in this paper are the infinitely smooth bump function

(5.3)
$$\psi_0(r) = \exp\left(-\frac{1}{(1-r^2)_+}\right),$$

for the geometry reconstruction, and the C^2 Wendland function [40]

(5.4)
$$\psi_0(r) = (4r+1)(1-r)_+^4$$

for the PDE solutions.

To construct the weight functions from the generating functions, we use Shepard's method [35]. We first define the sum

(5.5)
$$s(\underline{x}) = \sum_{j=1}^{P} \psi_j(\underline{x}),$$

which is positive for all $\underline{x} \in \bigcup_{j=1}^{P} \Omega_j$. Even though the sum is over all patches, it never has more than K non-zero terms at any given point due to the overlap condition. The weight function w_j supported on the patch Ω_j is given by

(5.6)
$$w_j(\underline{x}) = \frac{\psi_j(\underline{x})}{s(\underline{x})}.$$

Using a recursive differentiation rule for rational functions, we furthermore have

(5.7)
$$\nabla w_j = \frac{\nabla \psi_j - w_j \nabla s}{s},$$
$$\Delta w_j = \frac{\Delta \psi_j - 2\nabla w_j \cdot \nabla s - w \Delta s}{s},$$
$$\mathcal{H}_{w_j} = \frac{\mathcal{H}_{\psi_j} - \nabla w_j^{\mathsf{T}} \nabla s - \nabla s^{\mathsf{T}} \nabla w_j - w \mathcal{H}_s}{s}.$$

An example with two weight functions is shown in Figure 5.1. Within the regions where there is no overlap, the local weight function is identically 1 and all derivatives are zero. The formulas above are therfore only needed in the overlap regions. At the outer boundary of $\bigcup_{i=1}^{P} \Omega_j$, at the points where two (or more) patches intersect, the weight functions are discontinuous, since the effective overlap at these points is 0. In the interior the weight functions have the full smoothness of the generating function.



FIG. 5.1. Contour plots of the weight functions of two adjacent 2D patches.

6. The smooth implicit geometry representation. We represent the geometry of the thin structures we are studying implicitly, through a level set function $\ell(\underline{x})$ that is zero at the object surface. The gradient of the level set function at the object surface provides the normal direction, but we only require $\ell(\underline{x})$ to approximately represent the distance to the object surface with negative values inside and positive outside.

To construct $\ell(\underline{x})$ we use an RBF-PUM approximation of form (3.1)

(6.1)
$$\ell(\underline{x}) = \sum_{j=1}^{P} w_j(\underline{x})\ell_j(\underline{x}) = \sum_{j=1}^{P} w_j(\underline{x})\gamma_j(T_j(\underline{x}))$$

where γ_j are functions of form (3.3)–(3.4) defined over the reference patch. The least squares approximation problems for the functions $\gamma_j(\underline{x})$ are solved independently. By using the infinitely smooth multiquadric RBF for the local approximation and the infinitely smooth weight function (5.3), the geometry reconstruction also becomes infinitely smooth. We need to choose the number of basis functions n in the reference patch such that we ensure oversampling in all patches. The m_j data points in patch jgenerate $2m_j$ conditions based on the locations and normal directions. This provides the bound

$$(6.2) n \le \min_{j} 2m_j.$$

We are in fact adding some extra conditions as described below, but we keep this bound since it is the limit when considering only the data.

The local subproblems contain three types of conditions. The conditions are formulated in the physical domain, but implemented in the reference patch, where all local problems are solved. The first condition is the zero value at the object surface

(6.3)
$$\ell_j(\underline{y}_i^j) = \gamma_j(T_j(\underline{y}_i^j)) = 0, \quad \underline{y}_i^j \in Y_j^\rho,$$

where Y_j^{ρ} is an extension of the local data set Y_j with the data points of distance at most ρh outside the patch. By adding the extra conditions outside we improve the approximation at the intersection of the radial patch boundary and the object.

The second set of conditions guides the gradient of the level set function. We follow the approach in [30] and require the gradient to have a unit component in the initial normal direction. However, since we are working with local patch approximations and not stencils, we enforce these conditions at all data points instead of only at the stencil center as was done in [30]. We use (4.6) to transform the gradient to get

(6.4)
$$\underline{n}_{i}^{j} \nabla \ell_{j}(\underline{y}_{i}^{j})^{\mathsf{T}} = \underline{n}_{i}^{j} QS \nabla' \gamma_{j}(T_{j}(\underline{y}_{i}^{j}))^{\mathsf{T}} = 1, \quad \underline{y}_{i}^{j} \in Y_{j}^{\rho},$$

where \underline{n}_{i}^{j} is the initial normal at \underline{y}_{i}^{j} .

The conditions (6.3) and (6.4) are all positioned at the object surface. This implies that in regions of the patch that are far from the object, the level set function may misbehave, and for example introduce spurious zero level set curves. To stabilize the approximation, we add a third set of conditions. We start from the set X' of RBF centers in the reference patch. We compute the average node distance $h' = \frac{1}{n} \sum_{i=1}^{n} \min_{j \neq i} ||x'_i - x'_j||$ (which we use as an approximation of the fill distance) and use this to construct a node set X'_b with approximate node distance h' at the reference patch surface. We combine the boundary nodes with centers that are at least a distance $\alpha h'$ away from the object to get the set $X'_{j,e} = X'_b \cup \left\{ x'_k \in X' \mid \min_i ||x'_k - T_j(\underline{y}_i^j)|| > \alpha h' \right\}$. Then, for each $\underline{y}_i^{j,e}$ such that $T_j(\underline{y}_i^{j,e}) \in X'_{j,e}$ we find the index of the closest data point on the object $\pi_{i,j} = \arg\min_k ||\underline{y}_i^{j,e} - \underline{y}_k||$, where $\underline{y}_k \in Y$. We use the approximate distance from the boundary data in the normal direction as a condition for the value of the level set function at that point

(6.5)
$$\ell_j\left(\underline{y}_i^{j,e}\right) = \gamma_j(T_j(\underline{y}_i^{j,e})) = \left(\underline{y}_i^{j,e} - \underline{y}_{\pi_{i,j}}\right) \cdot \underline{n}_{\pi_{i,j}}, \quad T_j(\underline{y}_i^{j,e}) \in X'_{j,e}.$$

The three types of conditions are illustrated for a two-dimensional patch in Figure 6.1. Equations (6.3), (6.4), and (6.5) form the local least-squares problems. Solving these



FIG. 6.1. The local least-squares problem for one patch in the approximation of the 2D diaphragm data. The center points (×), the data points (colored dot and arrow), and the additional points $X_{j,e}$ (colored dot) are shown. The color corresponds to the function value and the arrows are the normal directions used in conditions (6.4) and (6.5).

gives the vectors \underline{u}_j of patch nodal values, which allows us to evaluate the local RBF approximation (3.3). Using (3.2) and (3.1) then provides the global level set function at all points within the union of patches.

The algorithm that we have derived has a number of parameters that we discuss qualitatively here, and more quantitatively in section 7. The number of patches Ptogether with the relative overlap δ_0 determine the number of data points m_j within a patch. Furthermore, δ_0 governs the transition between the solutions in adjacent patches. A small δ_0 can lead to large gradients in the transition regions, while a large δ_0 increases the computational cost. The number of center points n in the reference patch and the overlap δ_0 , together with the number of patches P, determine the oversampling, and accordingly how closely the local data is fitted.

The reference patch dimensions R_0 and H_0 are coupled parameters. Therefore, we use a scale parameter R_s , such that $R_0 = R_s \bar{R}$ and an anisotropy parameter H_s , such that $H_0 = H_s R_s \bar{H}$. When performing numerical experiments, we fix the scale parameter, and only vary the anisotropy parameter.

The multiquadric RBF $\phi(r) = \sqrt{1+r^2}$ can be equipped with a shape parameter ε by using $\phi(\varepsilon r)$ as the basis functions. A small shape parameter is appropriate when approximating a function that is smooth in the sense that it mainly/only contains low frequency Fourier modes. At the same time, the conditioning of the local problem grows with an inverse power of ε [33, 19]. The ill-conditioning issues can be removed by using a stable evaluation method [15, 14, 13, 11, 20, 43, 18]. However, here we choose to use the multiqadric RBF with an intermediate shape parameter value that gives a reasonable condition number. This makes the local approximations robust and well-behaved also away from the object surface.

The parameters ρ and α that determine which extra conditions to use outside and inside the domain are less crucial for the results, but the existence of extra conditions is important for the robustness of the local problem.

The effect of the most important parameters are shown in Figure 6.2. We see that increasing n leads to a closer fit of the data, which in this case means fitting the noise. When $H_s < 1$, it becomes hard to differentiate between the two surfaces, and the algorithm makes a mistake at the lower right part and creates a disconnected zero level set. When we remove the extra conditions inside the patch and use a large n the resulting level set function has spurious zero level sets both inside and outside the actual geometry. This effect is almost invisible for n = 21, but worsens with increasing

n. In these experiments, the local node set X' consisted of Halton nodes [16].



FIG. 6.2. Level set functions for the 2D diaphragm for different parameter choices. The zero level set, i.e., the surface of the reconstructed geometry is the black curve. The result for $H_0 = 1$, $R_0 = 1$, n = 21, P = 12, $\delta_0 = 0.25$, $\varepsilon = 1$, $\rho = 0.25$, $\alpha = 0.5$ is shown at the back. From back to front, for the other cases, the differences are for the first case n = 91, for the second case $H_0 = 0.25$, and for the third case n = 91 and $X_{j,e} = \emptyset$.

6.1. Geometry quality measures. To evaluate the reconstructed geometries numerically, we introduce two different quality measures. As a measure of smoothness, we use the mean Frobenius norm of the surface Hessian of the geometry, which is computed using the following steps:

- 1. The data points $\{\underline{y}_k\}_{k=1}^M$ are iteratively moved onto the geometry surface using the values and gradients of the levelset function $\ell(\underline{y})$.
- 2. The Hessian $\mathcal{H}_{\ell}(\underline{y})$ is computed at each surface point $\{\underline{y}_{k}^{\overline{S}}\}_{k=1}^{M}$ using (4.6).
- 3. For each surface point, we find a rotation Q_y , such that the normal at y becomes the dth unit vector. The surface Hessian is then given by

$$\mathcal{H}_{\ell}^{S}(\underline{y}) = \left(Q_{y}\mathcal{H}_{\ell}(\underline{y})Q_{y}^{\mathsf{T}}\right)_{1:d-1,1:d-1}$$

4. Finally, the mean Frobenius norm, F_H is computed as

$$F_H = \frac{1}{M} \sum_{k=1}^M \|\mathcal{H}_\ell^S(\underline{y}_k^S)\|_F$$

When least-squares approximation is employed for smoothing, the resulting smoothness is coupled to the amount of oversampling. That is, we expect F_H to decrease if we decrease the number of degrees of freedom. However, there is also a problem dependent lower bound for how many degrees of freedom we need, to capture the true structure behind the noisy data. The second quality measure is an indicator specific to thin structures, which we have found to be effective in detecting when issues related to lack of resolution appear. It requires the nodes to be annotated with inner, outer and edge, which can be done automatically. We let $K_{\rm in}$ be the index set corresponding to data points on the inner surface, and we let $\tau(\underline{y}_k)$ be the distance to the closest data point on the outer surface. We measure the relative loss of thickness T_L as

(6.6)
$$T_L = 1 - \min_{k \in K_{\rm in}} \frac{\tau(\underline{y}_k^S)}{\tau(\underline{y}_k)}.$$

The loss is not expected to become zero, but in a well-resolved situation it is on the order of the noise size.

7. Numerical results for the geometry reconstruction. The implementation is written in MATLAB and all experiments have been performed on standard laptop computers.

In the experiments we have chosen to use Halton nodes [16] for the RBF center points X' in the reference patch. The main reason is that low discrepancy points are insensitive to scaling and are therefore well-distributed both in the reference patch and in the physical domain. In the experiments, we have fixed R_s such that $R_0 = 1$ for P = 10 in 2D, and for P = 20 in 3D.

The set of parameters we can use to optimize the quality of the reconstructed geometry are listed and described in Table 7.1

TABLE 7.1The tunable parameters and their descriptions

0	Determines the minimum distance from the surface to the setue points in (6.5)							
	Determines the minimum distance from the surface to the extra points in (0.5) .							
	Increasing α will slightly decrease the oversampling factor.							
ρ	decides how far outside the patch boundaries data is collected for the local prob-							
	lems to form the data conditions (6.3) and (6.4). Increasing ρ increases the							
	oversampling slightly, but the main and desired effect is to stabilize the behavior							
	of the level set function where the patch intersects the geometry.							
H_s	Controls the aspect ratio of the reference patch. By increasing H_s we make the							
	contribution to the residual from the distance normal to the surface larger than							
	the contribution from the distance tangential to the surface. That is, we increase							
	the amount of surface smoothing. Also, the distance between the two surfaces							
	as well as the end curvature is decreased in the computational domain.							
n	This is the number of RBF center points in each patch, which has the strongest							
	influence on the total number of degrees of freedom.							
P	This is the number of patches, which is the second most important parameter in							
	terms of degrees of freedom.							
δ_0	is the relative overlap between patches. When $\delta_0 = 1$ patch neighbors on opposite							
	sides of one patch meet in the middle. Increasing δ_0 increases the oversampling,							
	since more data points are included in the local subproblems, while the number							
	of RBF center points in the patch is the same. The resolution of the geometry is							
	however decreased, since the same number of basis functions cover a larger area							
	in the physical domain.							
ε	The RBF shape parameter affects the surface Frobenius norm. With a larger							
	shape parameter more variations in the data, including the noise, are picked up.							

7.1. Results for 2D cross sections. For the 2D cross section data shown in Figure 2.1 one parameter at a time was varied to study its behavior for all combinations of n = 21, 28, and 36 and P = 10, 20, 30, and 40. In this phase, we mainly focus on the smoothness measure F_H .

The smoothness of the reconstruction is relatively insensitive to the distance ρh at which we pick up data outside the patch. The choice $\rho = 0.5$ performs well in all cases, and is significantly better than $\rho < 0.5$, especially for larger numbers of patches. We

keep this parameter fixed in all subsequent experiments. The reconstruction is even less sensitive to the distance αh between the object and the extra conditions inside the patch. We fix this parameter to $\alpha = 1.5$. For the overlap, the smoothness increases significantly for the range from $\delta_0 = 0.1$ to 0.9. However it is worth noticing that the loss in thickness increases with δ_0 . Depending on the data, different overlaps may be optimal. In the later experiments, we use both $\delta_0 = 0.5$ and $\delta_0 = 0.9$. The curvature measure F_H decreases with the shape parameter down to $\varepsilon \approx 0.2$ for the tested cases. For small values the thickness loss increases and the computations gradually become ill-conditioned, since we are solving the local systems directly. To stay on the robust side, we use $\varepsilon = 0.5$ in the subsequent experiments. The patch scaling is fixed to $R_s \approx 0.0215$, which makes $R_0 = 1$ for P = 10 patches. An increase of the anisotropy factor H_s leads to a smoother result, especially for larger P, whereas loss of thickness occurs to a larger degree for small H_s . In the later experiments, we use both $H_s = 1$ and $H_s = 2$.

To find reconstruction parameters that minimize the quality measures we optimized $F_H + T_L$ over P_0 for different n and the four combinations of overlap and aspect ratios. Since T_L is an order of magnitude larger than F_H , the priority becomes to find a reconstruction that keeps an even width throughout. The resulting parameters are reported in Table 7.2. We note that all choices of n gave good reconstruction results. If the even smaller n = 6 is used, we also get a good result in terms of quality, but points move more with respect to the given data.

TABLE 7.2 The best results for different n if $F_H + T_L$ is minimized over P_0 for the four combinations of

 δ_0 and H_s .

n	P_0	P	δ_0	H_s	F_H	T_L
10	30	34	0.5	1	0.024	0.20
15	- 33	37	0.5	1	0.022	0.14
21	22	24	0.5	1	0.016	0.11
28	12	12	0.5	1	0.013	0.13
36	11	11	0.5	1	0.014	0.12
45	11	11	0.5	1	0.014	0.12

45 11 11 0.5 1 0.014 0.12 The best result according to the objective function is for n = 21. The geometry and the quality measures are illustrated in Figure 7.1. In the left side of the domain, the loss of thickness seems to be correlated with the patch overlap regions, while this is less clear in the rest of the domain. In the right part of the figure, we can observe that if we allowed for a higher thickness loss, we could reach an overall smoother

result. To show that the reconstruction algorithm can handle different geometries, we run the same code with the same objective function for two other cross sections with n = 21. The results including best parameters and quality measures are shown in Figure 7.2.

7.2. Results for 3D reconstruction. We reconstruct 3D diaphragm geometries from two different patients. The reconstructions are shown in Figure 7.3 and Figure 7.4. Similarly to the reconstruction of the 2D cross sections we optimize for each parameter separately while fixing all others. We choose parameters that minimize the quality measure $F_H + T_L$ and all parameter choices are made using the geometry shown in Figure 7.3. The same parameters are used to reconstruct the more complex geometry.

The reconstruction is quite robust in terms of the distance for which we pick points



FIG. 7.1. The best reconstruction for n = 21 and P = 24 with $F_H = 0.016$ and $L_T = 0.11$ is shown together with the spatial distribution of the quality measures to the left. The black curve is the zero level set. Negative losses are set to zero and correspond to dark blue. The quality measures are shown as a function of the number of patches P to the right.



FIG. 7.2. Reconstructed diaphragm cross sections for n = 21 in two other planes. The best parameters for the left case are $P_0 = 40$, P = 46, $\delta_0 = 0.9$, and $H_s = 1$, leading to $F_H = 0.008$ and $T_L = 0.22$. The best parameters for the right case with two loops are $P_0 = 21$, P = 23, $\delta_0 = 0.5$, and $H_s = 1$, leading to $F_H = 0.013$ and $T_L = 0.14$.

outside the patch, ρh and the distance αh between the object and the extra conditions inside the patch. Here, the quality measure is minimized when $\alpha = 0$ and $\rho = 0$. We additionally find that the quality measure is fairly constant with respect to the chosen anisotropy ratio H_s of the reference patch, with slightly better results provided when $H_s = 1$. Moreover, similar to the 2D reconstruction, an increase in relative overlap δ_0 results in a significant increase in smoothness. However, it also results in a loss of thickness and the optimal value based on our metric is $\delta_0 = 0.4$. The optimal shape



FIG. 7.3. The best 3D diaphragm reconstruction, using n = 120 and P = 393 leading to $F_H = 0.053$ and $T_L = 0.18$ (left). The quality measures are shown as a function of the number of patches, with n = 120 (right).

parameter ε with respect to the resulting smoothness is relatively small ($\varepsilon = 0.25$), however the loss of thickness is at a maximum ($T_L = 1$). This means that some data points originally annotated with being on the inner surface have moved to the outer surface or vise versa. The optimal quality measure is obtained when $\varepsilon = 0.9$.

We note that point flipping between inner and outer surfaces is unique to the 3D reconstruction problems, were the distance between the inner and outer surface data points, as well as the inner and outer surface reconstruction, can be very small in specific areas. This can indicate different issues, either relating to the annotation of points, the data quality, or the reconstruction quality. Hence, in cases when this happens at isolated points, which indicates bad data quality or annotation, we do not include these points in the thickness loss measure T_L . However, in cases where multiple points in the same neighborhood exhibit the same issue, they cannot be ignored.

With the aforementioned parameters fixed, we vary the degrees of freedom to see how the quality measures are affected. We find the optimal number of patches P_0 for each choice of RBF center points n. Results are provided in Table 7.3.

TABLE 7.3 The best results for different n if $F_H + T_L$ is minimized over P_0 for fixed values of $\alpha = 0$, $\rho = 0$, $\delta_0 = 0.4$, $H_s = 1$ and $\varepsilon = 0.9$.

n	P_0	P	F_H	T_L
35	300	393	0.031	0.32
56	420	507	0.039	0.28
84	300	393	0.044	0.21
120	300	393	0.053	0.18
165	180	227	0.046	0.21
220	180	227	0.050	0.19

We obtain the optimal result according to the objective function when n = 120. The two reconstructed diaphragm geometries are shown in Figure 7.3 and Figure 7.4. In Figure 7.3 it is also shown how the two quality measures change when the patch number is increased. As is expected from the 2D examples, an increase in the number of patches P results in an increase in the Frobenius norm F_H , but a decrease in thickness loss T_L . For the more complex diaphragm geometry we also provide the smoothness measure over the surface. Note the high F_H values around the holes, where the vena cava and esophagus pass through the diaphragm, which make this a more challenging geometry to reconstruct.



FIG. 7.4. The second 3D diaphragm geometry reconstruction for n = 120 and P = 393 with $F_H = 0.067$ and $T_L = 0.72$ (left). The Frobenius norm plotted on the reconstructed surface of the diaphragm (left).

8. The PDE solution. Given the smooth geometry reconstruction, we want to use it for solving a PDE problem. In particular, we eventually want to model the biomechanical action of the diaphragm during respiration, see also [2, 38]. The base method, RBF-PUM, is the same as for the reconstruction, but there are some differences which we describe in detail in this section. In this paper we solve a Poisson problem with Dirichlet boundary conditions as a proof of concept. We manufacture the right hand side functions based on an infinitely smooth solution function u and write the problem as

(8.1)
$$\Delta u(\underline{x}) = f_i(\underline{x}), \quad \underline{x} \in \Omega,$$

(8.2)
$$u(\underline{x}) = f_b(\underline{x}), \quad \underline{x} \in \partial \Omega.$$

We sample the PDE and boundary condition at the point sets Y_i and Y_b , respectively. To form the global linear system, we use (3.6) with $\mathcal{L}_k = \Delta$ and $f = f_i$ for $\underline{y}_k \in Y_i$ and $\mathcal{L}_k = I$ and $f = f_b$ for $\underline{y}_k \in Y_b$. For the derivatives of the local approximations we use (4.6) and for the derivatives of the weight function we also use (5.7). The scaling of the equations is implemented as described in section 3, using the power $h^{-1/2}$ for the Dirichlet conditions.

To generate a cover for the PDE computation, we need a point set on the computed object surface. We use the initial data point cloud from the reconstruction problem, and move these points iteratively onto the surface using the gradient of the level set function. This works well when the level set function is free of spurious zero level set curves. Then we run the cover algorithm derived in section 4.

To form the interior node set Y_i , we first generate Halton nodes within each patch. Then we use the level set function to determine which points are inside and which are outside, and finally we remove points that are close to each other. Least-squares RBF methods are not sensitive to the exact distribution of the sampling points as long as these are quasi uniform [23, 37, 21]. For the surface point set Y_b , we use the mesh representation of the initial point cloud. We refine or decimate the mesh and move all new points onto the surface. It would be possible to construct a surface node generation algorithm that uses the level set function directly, but we have not investigated this aspect here.

Using RBF-PUM, there are two standard modes of refinement. Patch refinement with a fixed number of nodes per patch leads to algebraic convergence and node refinement with fixed patches leads to spectral convergence for infinitely smooth RBFs [23]. For both refinement modes, the shape parameter should be kept fixed. If a stable evaluation method is not used, the ill-conditioning of the local subproblems increases with the refinement level [23]. Therefore, we use Gaussian RBFs and the RBF-QR method [13, 20] for stable evaluation when solving PDEs. We note that to derive the transformed Laplacian section 4, we need the second derivatives with respect to each coordinate direction. These derivatives were not present in the 3D RBF-QR implementation derived in [20], but we have now upgraded the implementation with all second derivatives [22]. Furthermore, we replace the infinitely smooth bump function with the C^2 Wendland function (5.4) as generating function for the weights. This implies that we only require C^2 continuity of the global solution, while retaining the global high-order convergence arising from the high-order local approximations [41, 10, 23].

In the convergence tests performed here we use the patch refinement approach which is the most appropriate for large scale problems. It was shown numerically in [23] that least-squares RBF-PUM is stable under patch refinement. For our particular setting, we need to keep in mind that the computational domain is (several instances of) the reference patch. By keeping the scale factors R_s and H_s constant under refinement, we ensure that refining the physical domain corresponds to refining the computational domain. We choose to do the refinement only in the radial direction. That is, we keep the structure with one layer of patches. The node layout in the reference patch as well as the physical domain is quasi uniform, and as a practical measure of the worst case fill distance in the physical domain, we use

(8.3)
$$h = \max_{1 \le j \le P} \left(\frac{|\Omega_j|}{n}\right)^{\frac{1}{d}} = C_d \max_{1 \le j \le P} \left(\frac{H_j R_j^{d-1}}{n}\right)^{\frac{1}{d}}$$

Since H_j should be proportional to the width of the geometry and almost constant, the fill distance is approximately proportional to $R_j^{(d-1)/d}$. That is, by dividing a patch in half in 2D, we only reduce the fill distance by $1/\sqrt{2}$.

patch in half in 2D, we only reduce the fill distance by $1/\sqrt{2}$. Let $\mathcal{I}_h(u) = \sum_{j=1}^P w_j \mathcal{I}_h(u_j)$, where $u_j = u|_{\Omega_j}$, be the RBF-PUM interpolant of the function u, and note that $u(\underline{x}) = \sum_{j=1}^P w_j u_j$, $\underline{x} \in \Omega$ due to the partition of unity property. Then we can express the interpolation error as

(8.4)
$$\mathcal{E}_I = \mathcal{I}_h(u) - u = \sum_{j=1}^P w_j (I_h(u_j) - u_j) \equiv \sum_{j=1}^P w_j e_j, \quad \underline{x} \in \Omega.$$

Similarly, the consistency error for the Laplacian operator becomes [41, 23]

(8.5)
$$\mathcal{E}_{\Delta} = \Delta \sum_{j=1}^{P} w_j e_j = \sum_{j=1}^{P} (\Delta w_j e_j + 2\nabla w_j \cdot \nabla e_j + w_j \Delta e_j), \quad \underline{x} \in \Omega.$$

The derivatives of the C^k partition of unity weight functions are typically expressed in terms of the radius of the patch, especially in a setting with spherical patches [41, 23]. In our context with a non-uniform distribution of cylindrical patches, we need to be more careful. The largest size of the derivatives depends on the smallest absolute overlap δ between patches. We use the following estimate

(8.6)
$$\|D^{\beta}w\|_{\infty} = G_{|\beta|}\delta^{-\beta}, \quad |\beta| \le k.$$

For the local interpolation error, we use results derived in [31]. For our specific setting, we first make the estimate in the computational domain, and in a second step include the transformation. We let $v_j(x') = u_j(x)$ and define n_p as the dimension of a polynomial space of degree p in d dimensions. Then we use the following error estimate;

(8.7)
$$\|D^{\beta}(I_h(v_j) - v_j)\|_{\infty} \le C_{|\beta|} h^{p+1-|\beta|-\frac{d}{2}} \|v_j\|_{\mathcal{N}(\Omega_0)},$$

where $n_p \leq n < n_{p+1}$, the constant $C_{|\beta|}$ depends on the shape of the reference patch and on d, and $\|\cdot\|_{\mathcal{N}(\cdot)}$ denotes the native space of the infinitely smooth radial basis function that is used.

Remark 8.1. From the theoretical perspective, the results in [31] require h to be small enough for the estimates to be valid. In our setting h is related to n and cannot be decreased independently. The theoretical upper bound on h is restrictive, while in numerical experiments the optimal rate is achieved also for larger h [23].

We finally use (4.6) to provide estimates for e_j in the physical domain

(8.8)
$$\|\nabla e_j\|_{\infty} \leq \|\nabla'(I_h(v_j) - v_j)S^{\mathsf{T}}Q^{\mathsf{T}}\|_{\infty} \leq d \max_{1 \leq i \leq d} s_{ii}C_1h^{p-\frac{d}{2}}\|v_j\|_{\mathcal{N}(\Omega_0)},$$

(8.9)
$$\|\Delta e_j\|_{\infty} \leq \sum_{i=1}^d s_{ii}^2 \|\partial^2 / \partial x_i'^2 (I_h(v_j) - v_j)\|_{\infty} \leq d \max_{1 \leq i \leq d} s_{ii} C_2 h^{p-1-\frac{d}{2}} \|v_j\|_{\mathcal{N}(\Omega_0)}.$$

Combining (8.4)–(8.9) we get the consistency estimates

(8.10)
$$\|\mathcal{E}_I\|_{\infty} \le KG_0 C_0 h^{p+1-\frac{d}{2}} \max_j \|v_j\|_{\mathcal{N}(\Omega_0)},$$

(8.11)
$$\begin{aligned} \|\mathcal{E}_{\Delta}\|_{\infty} &\leq K \left(dG_{2} \delta^{-2} C_{0} h^{p+1-\frac{d}{2}} + 2dG_{1} \delta^{-1} d \max_{i} s_{ii} C_{1} h^{p-\frac{d}{2}} + G_{0} d \max_{i} s_{ii} C_{2} h^{p-1-\frac{d}{2}} \right) \max_{j} \|v_{j}\|_{\mathcal{N}(\Omega_{0})}, \end{aligned}$$

where K is the maximum number of patches overlapping at any given point. The best convergence rate that can be achieved is $\mathcal{O}(h^{p-1-\frac{d}{2}})$. This requires the absolute overlap δ to go to zero with the same rate as h or slower than h. We see from (8.3) that $h \propto R_j^{\frac{d-1}{d}}$, which with $\delta = \delta_0 R_j$ means that the relative overlap δ_0 must increase as $1/\sqrt[d]{R_j}$ during refinement to achieve the optimal rate. This is an effect of the refinement mode we have chosen where patches are split only in d-1 directions. When patches are split uniformly, keeping δ_0 constant is instead optimal. The scaling of the Dirichlet conditions with $h^{-1/2}$ is small enough to not affect the optimal rate.

In the convergence experiments, the relative maximum norm error is approximated as

(8.12)
$$E = \frac{\|\tilde{u} - u\|_{\ell_{\infty}(Y)}}{\|u\|_{\ell_{\infty}(Y)}}.$$

9. Numerical results for the PDE solution. As the first test problem, we solve the Poisson problem (8.1)–(8.2) in the 2D cross section of the diaphragm, with right hand side data generated from the solution function

(9.1)
$$u_{2D} = \frac{\sin(24(\tilde{x}_1 - 0.1)^2)\cos((\tilde{x}_2 - 0.3)^2) + \sin^2(6(\tilde{x}_2 - 0.5)^2)}{1 + 2\tilde{x}_1^2 + \tilde{x}_2^2},$$

and in the 3D geometry with right hand side data generated from

(9.2)
$$u_{3D} = \frac{\sin(4\pi\tilde{x}_1\tilde{x}_2\tilde{x}_3) + \cos(4\pi(\tilde{x}_2 - 0.3)^2)}{1 + 2\tilde{x}_1^2 + \tilde{x}_2^2 + 0.5\tilde{x}_3^2}$$

These functions provide variation both along and across the diaphragm. The scaled variables are given by $\tilde{x}_i = x_i / \max_{x \in \Omega} |x_i|$. The solution functions are illustrated in Figure 9.1.



FIG. 9.1. The manufactured solution functions u_{2D} (9.1) and u_{3D} (9.2). For the latter, bright yellow corresponds to +1 and dark blue to -1.

In the 2D experiments, which run up to a large number of patches, we have not increased the relative overlap as fast as suggested by the theory. The reason is that if the overlap becomes too large, patches that are not neighbours start to overlap, and this is an overlap that we cannot control, and which can produce large gradients in the weight functions. Instead of using $\delta_0 \propto \bar{R}^{-1/d}$, we have used $\delta_0 \propto \bar{R}^{-\frac{1}{2d}}$ for this case.

The purpose of the 2D experiment is to investigate how the result depends on the method parameters and how sensitive the parameter choices are. The left subplot of Figure 9.2 shows convergence as a function of the fill distance and the right one shows error as a function of time. The best theoretical convergence rates from (8.11)



FIG. 9.2. The relative maximum error for the 2D problem is shown for different numbers of local points n as a function of the fill distance h (left subplot) and as a function of the run time in seconds (right subplot). The parameters not indicated in the subplots are chosen as the default values q = 3, $\varepsilon = 1$, $H_s = 1$, and $\delta_0 = 0.1$.

become 3, 5, 7, and 9 for these n. The numerical results seem to be around one order

20

better. We note that increasing the number of local points allows us to reach smaller errors for the same fill distance. To reach a smaller error in the most efficient way also requires gradually increasing n.

To investigate the parameter dependence, we have solved the PDE problem for all combinations of n = 21, 28, 36, 45, 55, 78, and 91 and all 24 values of P used for the plots for different parameter combinations. A measure of goodness was defined by looking for results that are more accurate and/or faster than the average performance when the other parameters are equal. Then for a selection of parameter values we compute in which percentage of cases each choice was the best and in what percentage of cases it failed, where fail means that the result is at least 5 orders of magnitude less accurate than for other parameter values. Note that a smaller shape parameter

TABLE 9.1

The percentage of wins $B(\cdot)$, and fails $F(\cdot)$ for each tested parameter value. The shaded row indicates the default values used in the experiments.

ε	$B(\varepsilon)$	H_s	$B(H_s)$	δ_0	$B(\delta_0)$	$F(\delta_0)$	q	B(q)	F(q)
0.50	27						2.0	51	5
0.75	39	0.75	37	0.05	23	1	2.5	31	0
1.00	34	1	51	0.1	61	0	3.0	18	0
		1.25	12	0.2	16	0			

value than 1 performs a little bit better, but the trend of the convergence curve was a bit more erratic, and we therefore picked $\varepsilon = 1$ for improved stability. A smaller oversampling factor q improves the efficiency, but the method can fail, so also here, we picked q = 3 for stability. Also for the overlap, a small overlap is efficient, but when too small, the method can fail. It should also be noted that for all stable choices of parameters, we got a similar convergence and performance. That is, the main sensitivities can be understood and for the rest, the penalty for making a suboptimal choice is not too large.

The error for two different numbers of patches is illustrated in Figure 9.3. Especially in the left subfigure, it is clear that the error is a bit larger in the overlap regions, where the weight functions have large gradients.



FIG. 9.3. The logarithm in base 10 of the error in the RBF solutions for n = 78 and P = 25 (left) and P = 69 (right).

Convergence and efficiency for the 3D problem are shown in Figure 9.4. The theoretical convergence rates from (8.11) here become 0.5, 2.5, 4.5, and 6.5. Note that here the amount of RAM in the laptop limits how large problems can be solved. Therefore the curves for larger n have fewer points. The values of P_0 that were used are 40, 60:10:90, 100:20:300, 360, 400. The numerical convergence rates are also here better than expected, but in a less systematic way. We would need to run a larger range of problem sizes to really see the trends. The behavior with respect to efficiency is similar to that for the 2D case. The number of local points n should increase when the error tolerance is lowered.



FIG. 9.4. The relative maximum error for the 3D problem is shown for different numbers of local points n as a function of the fill distance h (left subplot) and as a function of the run time in seconds (right subplot) using q = 5, $\varepsilon = 1.2$, $H_s = 1$, and $\delta_0 = 0.1$.

The errors in the 3D solution for one set of parameters is shown in different slices in Figure 9.5. As can be seen in the left subplot, the patch structure in 3D becomes non-trivial, making it hard to say if and how the error is affected by the patch overlaps. However, the error distribution appears to be relatively even within each slice.



FIG. 9.5. The logarithm in base 10 of the errors in one slice (left) and in ten slices (right) in the RBF-PUM solutions to the 3D problem for n = 220 and P = 100. In the left subfigure, also the intersections of the plane with the cover is shown.

10. Conclusions. In this paper, we have developed an geometry adaptive least squares RBF-PUM. We first used the method for reconstruction of a geometry from a noisy point cloud. The reconstruction is stable enough that it can produce good results for different parameter choices, and it does not produce spurious zero level sets. We showed how the parameters can be optimized using quality measures. This approach can be used for different types of thin geometries, but a similar method could also be defined for general surfaces, where the patch layer would then follow the surface, not the geometry. In the second part of the paper we showed that we can achieve a sustained and predictable convergence behaviour in our reconstructed 2D and 3D geometries, in spite of the irregular patch and node layouts.

REFERENCES

- S. AROUDJ, P. SEEMANN, F. LANGGUTH, S. GUTHE, AND M. GOESELE, Visibility-consistent thin surface reconstruction using multi-scale kernels, ACM Trans. Graph., 36 (2017), 187 (13 pages), https://doi.org/10.1145/3130800.3130851.
- [2] N. CACCIANI, E. LARSSON, A. LAURO, M. MEGGIOLARO, A. SCATTO, I. TOMINEC, AND P.-F. VILLARD, A first meshless approach to simulation of the elastic behaviour of the diaphragm, in Spectral and High Order Methods for Partial Differential Equations ICOSAHOM 2018, S. J. Sherwin, D. Moxey, J. Peiró, P. E. Vincent, and C. Schwab, eds., vol. 134 of Lect. Notes Comput. Sci. Eng., Springer International Publishing, Cham, 2020, pp. 501–512, https://doi.org/10.1007/978-3-030-39647-3_40.
- [3] J. C. CARR, R. K. BEATSON, J. B. CHERRIE, T. J. MITCHELL, W. R. FRIGHT, B. C. MC-CALLUM, AND T. R. EVANS, *Reconstruction and representation of 3D objects with radial basis functions*, in SIGGRAPH '01: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, New York, NY, USA, 2001, Association for Computing Machinery, pp. 67–76, https://doi.org/10.1145/383259.383266.
- [4] J. C. CARR, R. K. BEATSON, B. C. MCCALLUM, W. R. FRIGHT, T. J. MCLENNAN, AND T. J. MITCHELL, Smooth surface reconstruction from noisy range data, in Proceedings of the 1st International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia, GRAPHITE '03, New York, NY, USA, 2003, Association for Computing Machinery, pp. 119–126, https://doi.org/10.1145/604471.604495.
- [5] G. CASCIOLA, D. LAZZARO, L. MONTEFUSCO, AND S. MORIGI, Shape preserving surface reconstruction using locally anisotropic radial basis function interpolants, Comput. Math. Appl., 51 (2006), pp. 1185–1198, https://doi.org/10.1016/j.camwa.2006.04.002.
- [6] R. CAVORETTO AND A. DE ROSSI, Error indicators and refinement strategies for solving Poisson problems through a RBF partition of unity collocation scheme, Appl. Math. Comput., 369 (2020), pp. 124824, 18, https://doi.org/10.1016/j.amc.2019.124824, https://doi.org/ 10.1016/j.amc.2019.124824.
- [7] R. CAVORETTO, A. DE ROSSI, G. E. FASSHAUER, M. J. MCCOURT, AND E. PERRACCHIONE, Anisotropic weights for RBF-PU interpolation with subdomains of variable shapes, in Numerical mathematics and advanced applications—ENUMATH 2017, F. A. Radu, K. Kumar, I. Berre, J. M. Nordbotten, and I. S. Pop, eds., vol. 126 of Lect. Notes Comput. Sci. Eng., Springer, Cham, 2019, pp. 93–101, https://doi.org/10.1007/978-3-319-96415-7_6.
- [8] G. CHENG AND V. SHCHERBAKOV, Anisotropic radial basis function methods for continental size ice sheet simulations, J. Comput. Phys., 372 (2018), pp. 161–177, https://doi.org/10. 1016/j.jcp.2018.06.020, https://doi.org/10.1016/j.jcp.2018.06.020.
- K. P. DRAKE, E. J. FUSELIER, AND G. B. WRIGHT, Implicit surface reconstruction with a curl-free radial basis function partition of unity method, 2021, https://arxiv.org/abs/2101. 05940.
- [10] G. E. FASSHAUER, Meshfree approximation methods with MATLAB, vol. 6 of Interdisciplinary Mathematical Sciences, World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2007, https://doi.org/10.1142/6437, http://dx.doi.org/10.1142/6437.
- [11] G. E. FASSHAUER AND M. J. MCCOURT, Stable evaluation of Gaussian radial basis function interpolants, SIAM J. Sci. Comput., 34 (2012), pp. A737–A762, https://doi.org/10.1137/ 110824784, https://doi.org/10.1137/110824784.
- [12] B. FORNBERG AND N. FLYER, A primer on radial basis functions with applications to the geosciences, vol. 87 of CBMS-NSF Regional Conference Series in Applied Mathematics, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2015, https: //doi.org/10.1137/1.9781611974041.ch1.
- [13] B. FORNBERG, E. LARSSON, AND N. FLYER, Stable computations with Gaussian radial basis functions, SIAM J. Sci. Comput., 33 (2011), pp. 869–892, https://doi.org/10.1137/ 09076756X, https://doi.org/10.1137/09076756X.
- [14] B. FORNBERG AND C. PIRET, A stable algorithm for flat radial basis functions on a sphere, SIAM J. Sci. Comput., 30 (2007), pp. 60–80, https://doi.org/10.1137/060671991.
- [15] B. FORNBERG AND G. WRIGHT, Stable computation of multiquadric interpolants for all values of the shape parameter, Comput. Math. Appl., 48 (2004), pp. 853–867, https://doi.org/10. 1016/j.camwa.2003.08.010.
- [16] J. H. HALTON, On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals, Numer. Math., 2 (1960), pp. 84–90, https://doi.org/10.1007/ BF01386213, https://doi.org/10.1007/BF01386213.
- [17] INVIVE, The invive project, 2017, https://www.it.uu.se/research/scientific_computing/ project/rbf/biomech.
- [18] K. KORMANN, C. LASSER, AND A. YUROVA, Stable interpolation with isotropic and anisotropic gaussians using hermite generating function, SIAM J. Sci. Comput., 41 (2019), pp. A3839–

A3859, https://doi.org/10.1137/19M1262449, https://doi.org/10.1137/19M1262449.

- [19] E. LARSSON AND B. FORNBERG, Theoretical and computational aspects of multivariate interpolation with increasingly flat radial basis functions, Comput. Math. Appl., 49 (2005), pp. 103–130, https://doi.org/10.1016/j.camwa.2005.01.010.
- [20] E. LARSSON, E. LEHTO, A. HERYUDONO, AND B. FORNBERG, Stable computation of differentiation matrices and scattered node stencils based on Gaussian radial basis functions, SIAM J. Sci. Comput., 35 (2013), pp. A2096–A2119, https://doi.org/10.1137/120899108, https://doi.org/10.1137/120899108.
- [21] E. LARSSON, B. MAVRIČ, A. MICHAEL, AND F. POOLADI, A numerical investigation of some RBF-FD error estimates, Dolomites Res. Notes Approx., 15 (2022), pp. 78–95.
- [22] E. LARSSON AND A. MICHAEL, Computing mixed derivatives in 3D with the RBF-QR method. manuscript in preparation, 2024.
- [23] E. LARSSON, V. SHCHERBAKOV, AND A. HERYUDONO, A least squares radial basis function partition of unity method for solving PDEs, SIAM J. Sci. Comput., 39 (2017), pp. A2538– A2563, https://doi.org/10.1137/17M1118087.
- [24] B. LÉVY AND Y. LIU, L_p centroidal Voronoi tessellation and its applications, ACM Trans. Graphics, 29 (2010), 119 (11 pages), https://doi.org/10.1145/1778765.1778856.
- [25] S. LIU, C. C. WANG, G. BRUNNETT, AND J. WANG, A closed-form formulation of HRBF-based surface reconstruction by approximate solution, Comput.-Aided Des., 78 (2016), pp. 147– 157, https://doi.org/10.1016/j.cad.2016.05.001.
- [26] C. A. MICCHELLI, Interpolation of scattered data: distance matrices and conditionally positive definite functions, Constr. Approx., 2 (1986), pp. 11–22, https://doi.org/10.1007/ BF01893414.
- [27] Y. OHTAKE, A. BELYAEV, AND H.-P. SEIDEL, Sparse surface reconstruction with adaptive partition of unity and radial basis functions, Graph. Models, 68 (2006), pp. 15–24, https://doi.org/10.1016/j.gmod.2005.08.001.
- [28] R. PAN AND V. SKALA, A two-level approach to implicit surface modeling with compactly supported radial basis functions, Eng. Comput., 27 (2011), pp. 299–307, https://doi.org/ 10.1007/s00366-010-0199-1.
- [29] C. PIRET, The orthogonal gradients method: a radial basis functions method for solving partial differential equations on arbitrary surfaces, J. Comput. Phys., 231 (2012), pp. 4662–4675, https://doi.org/10.1016/j.jcp.2012.03.007.
- [30] C. PIRET AND J. DUNN, Fast RBF OGr for solving PDEs on arbitrary surfaces, in Numerical Computations: Theory and Algorithms (NUMTA-2016), Y. D. Sergeyev, D. E. Kvasov, F. Dell'Accio, and M. S. Mukhametzhanov, eds., vol. 1776 of AIP Conf. Proc., Amer. Inst. Phys., Melville, NY, 2016.
- [31] C. RIEGER AND B. ZWICKNAGL, Sampling inequalities for infinitely smooth functions, with applications to interpolation and machine learning, Adv. Comput. Math., 32 (2010), pp. 103–129, https://doi.org/10.1007/s10444-008-9089-0.
- [32] A. SAFDARI-VAIGHANI, A. HERYUDONO, AND E. LARSSON, A radial basis function partition of unity collocation method for convection-diffusion equations arising in financial applications, J. Sci. Comput., 64 (2015), pp. 341–367, https://doi.org/10.1007/s10915-014-9935-9, https://doi.org/10.1007/s10915-014-9935-9.
- [33] R. SCHABACK, Multivariate interpolation by polynomials and radial basis functions, Constr. Approx., 21 (2005), pp. 293–317, https://doi.org/10.1007/s00365-004-0585-2.
- [34] I. J. SCHOENBERG, Metric spaces and completely monotone functions, Ann. of Math. (2), 39 (1938), pp. 811–841, https://doi.org/10.2307/1968466.
- [35] D. SHEPARD, A two-dimensional interpolation function for irregularly-spaced data, in Proceedings of the 1968 23rd ACM National Conference, ACM '68, New York, NY, USA, 1968, ACM, pp. 517–524, https://doi.org/10.1145/800186.810616.
- [36] I. TOMINEC AND E. BREZNIK, An unfitted RBF-FD method in a least-squares setting for elliptic PDEs on complex geometries, J. Comput. Phys., 436 (2021), 110283 (24 pages), https: //doi.org/10.1016/j.jcp.2021.110283.
- [37] I. TOMINEC, E. LARSSON, AND A. HERYUDONO, A least squares radial basis function finite difference method with improved stability properties, SIAM J. Sci. Comput., 43 (2021), pp. A1441–A1471, https://doi.org/10.1137/20M1320079.
- [38] I. TOMINEC, P. VILLARD, E. LARSSON, V. BAYONA, AND N. CACCIANI, An unfitted radial basis function generated finite difference method applied to thoracic diaphragm simulations, 2021, https://arxiv.org/abs/2103.03673, https://arxiv.org/abs/2103.03673.
- [39] P.-F. VILLARD, P. BOSHIER, F. BELLO, AND D. GOULD, Virtual reality simulation of liver biopsy with a respiratory component, in Liver Biopsy, H. Takahashi, ed., InTechOpen, 2011, https://doi.org/10.5772/22033.

- [40] H. WENDLAND, Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree, Adv. Comput. Math., 4 (1995), pp. 389–396, https://doi.org/10.1007/ BF02123482.
- [41] H. WENDLAND, Fast evaluation of radial basis functions: methods based on partition of unity, in Approximation theory, X (St. Louis, MO, 2001), Innov. Appl. Math., Vanderbilt Univ. Press, Nashville, TN, 2002, pp. 473–483.
- [42] R. M. WHEBELL, T. J. MORONEY, I. W. TURNER, R. PETHIYAGODA, AND S. W. MCCUE, Implicit reconstructions of thin leaf surfaces from large, noisy point clouds, Appl. Math. Model., 98 (2021), pp. 416–434, https://doi.org/10.1016/j.apm.2021.05.014.
- [43] G. B. WRIGHT AND B. FORNBERG, Stable computations with flat radial basis functions using vector-valued rational approximations, J. Comput. Phys., 331 (2017), pp. 137–156, https: //doi.org/10.1016/j.jcp.2016.11.030, https://doi.org/10.1016/j.jcp.2016.11.030.
- [44] Q. XIA, M. Y. WANG, AND X. WU, Orthogonal least squares in partition of unity surface reconstruction with radial basis function, in Geometric Modeling and Imaging-New Trends (GMAI'06), 2006, pp. 28-33, https://doi.org/10.1109/GMAI.2006.40.
- [45] J. YANG, Z. WANG, C. ZHU, AND Q. PENG, Implicit surface reconstruction with radial basis functions, in Computer Vision and Computer Graphics. Theory and Applications, J. Braz, A. Ranchordas, H. J. Araújo, and J. M. Pereira, eds., Springer Berlin Heidelberg, 2008, pp. 5–12, https://doi.org/10.1007/978-3-540-89682-1_1.