IMPLICIT SURFACE RECONSTRUCTION WITH A CURL-FREE RADIAL BASIS FUNCTION PARTITION OF UNITY METHOD

KATHRYN P. DRAKE*, EDWARD J. FUSELIER[†], AND GRADY B. WRIGHT*

Abstract. Surface reconstruction from a set of scattered points, or a point cloud, has many applications ranging from computer graphics to remote sensing. We present a new method for this task that produces an implicit surface (zero-level set) approximation for an oriented point cloud using only information about (approximate) normals to the surface. The technique exploits the fundamental result from vector calculus that the normals to an implicit surface are curl-free. By using curl-free radial basis function (RBF) interpolation of the normals, we can extract a potential for the vector field whose zero-level surface approximates the point cloud. We use curl-free RBFs based on polyharmonic splines for this task, since they are free of any shape or support parameters. To make this technique efficient and able to better represent local sharp features, we combine it with a partition of unity (PU) method. The result is the curl-free partition of unity (CFPU) method. We show how CFPU can be adapted to enforce exact interpolation of a point cloud and can be regularized to handle noise in both the normals and the point positions. Numerical results are presented that demonstrate how the method converges for a known surface as the sampling density increases, how regularization handles noisy data, and how the method performs on various problems found in the literature.

Key words. Implicit surface reconstruction; Meshfree; Point clouds; Radial basis functions; Partition of unity; Curl-free; Potential; Polyharmonic splines; Smoothing splines

AMS subject classifications. 41-04, 41A15, 41A29, 65D07, 65D10, 65D15, 65D17, 68U07

1. Introduction. The problem of reconstructing a surface from a set of unorganized points, or a point cloud, has been used in a variety of applications, including computer graphics, computeraided design, medical imaging, image processing, manufacturing, and remote sensing. Many common methods developed to address this problem produce an implicit surface (zero-level set) using "oriented" point clouds, which involve the unstructured points as well as their corresponding normal vectors. In this paper, we present the Curl-free Radial Basis Function Partition of Unity (CFPU) method for this problem. The method is based on the following result from vector calculus: If $f: \mathbb{R}^d \longrightarrow \mathbb{R}$ defines a zero-level set \mathcal{P} (i.e. implicit curve for d=2 or surface for d=3) and \mathbf{n} is a normal vector to \mathcal{P} , then **n** is curl-free. This follows since **n** is proportional to ∇f and the curl of a gradient field is zero. Given a set of points $\{\mathbf{x}_1,\ldots,\mathbf{x}_N\}\subset\mathcal{P}$ equipped with oriented normal vectors $\{\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_N\}$, we seek to recover the potential f, such that $\nabla f \approx \mathbf{n}$ at every point \mathbf{x}_i . The method we use to recover f comes, in part, from a curl-free RBF approximant to the normal vectors. These approximants were introduced in [1, 15, 21] for generating an analytically curl-free approximant of curl-free field from scattered samples of the field. The key to our method lies in the feature that a scalar potential can be extracted from these vector approximants that can be used to approximate the implicit surface \mathcal{P} .

Implementing curl-free approximants globally is too computationally expensive, requiring the solution of a dN-by-dN system. To bypass this issue, we follow a similar approach to [17] and combine the technique with a partition of unity (PU) method. This allows the potential f to be solved for locally on patches involving $n \ll N$ points. These local potentials can then be blended together into a global potential whose zero-level set approximates \mathcal{P} . An added benefit of this approach is that it appears to be better equipped at recovering sharp features, which many global methods lack. The CFPU method can also be adapted to enforce exact interpolation of the surface and can be regularized to handle noisy data. Finally, by using curl-free polyharmonic splines (PHS) and curl-free vector polynomials to construct the local approximants, the method is free from shape and scaling parameters, which are common to other RBF methods and for which good values can

^{*}Department of Mathematics, Boise State University, Boise, ID (KathrynDrake@u.boisestate.edu, gradywright@boisestate.edu)

[†]Department of Mathematics, High Point University, High Point, NC (efuselie@highpoint.edu)

be computationally expensive to determine.

The rest of the paper is structured as follows. For the remainder of Section 1, we briefly discuss relevant surface reconstruction methods and compare them to CFPU. In Section 2, we provide background on curl-free RBF approximation and how it can be used for surface reconstruction. We then introduce the CFPU method in Section 3, along with modifications for exact surface interpolation and regularization. We discuss computational considerations and results of the CFPU method in Sections 4 and 5, respectively. Finally, we provide concluding remarks in Section 6.

1.1. Relationship to previous work. Reconstructing a surface from an unorganized point cloud has been extensively studied in literature since the 1990s. Some of the approaches involve signed distance methods [7, 24], RBF-based methods [9, 10, 19, 33–35, 37, 44, 47, 49, 51], partition of unity methods [23, 36, 38, 45], and methods which turn the reconstruction problem into a Poisson or screened Poisson problem [26, 27]. While a comprehensive review of the aforementioned methods is beyond the scope of this paper, we will briefly compare and contrast the methods related to CFPU; the interested reader is directed to the survey papers [5, 6] for a more comprehensive review.

The first closely related methods express the surface reconstruction problem as the solution to a (screened) Poisson equation [26, 27] (the so called indicator function approach [5]). Similar to CFPU, these methods rely on the fact that the normals of an oriented point cloud are the gradient of a potential. They use the divergence of the normals to get a (screened) Poisson equation that can be solved directly for the potential f. A weak formulation of the equation is used which replaces the divergence of the normals with integrals in the surfaces ambient space that involve an extension of the normals from the surface. One issue with these methods is that they require solving a global set of equations to recover f. In contrast, the CFPU method does not require extending the normals from the surface or computing their divergence since it reconstructs the potentials locally using curl-free RBF approximants of the normals directly from the surface. Furthermore, the method produces the global potential f by blending the local potentials together, which makes the method more amenable to parallel implementations.

A second closely related method to CFPU is HRBF Implicits [34]. The idea behind this method is to interpolate the potential (which is taken to be zero) at the point cloud samples and the normal vectors at these samples using a Hermite RBF interpolant. The surface \mathcal{P} is then determined directly from this interpolant. Our method, by comparison, only interpolates the normal vectors using a specially constructed matrix-valued kernel that allows us to extract a potential for the field directly, which can be used to approximate \mathcal{P} . This allows us to immediately reduce the size of the linear systems that need to be solved by 33% for 2D problems and 25% for 3D problems.

There are other surface reconstruction methods from the literature based on RBFs that are less closely related to the CFPU method, but are still relevant. Global RBF methods were initially used for modeling surfaces due to their ability to handle sparse point clouds; however, their global nature restricted their applications to small problems [10, 46]. Subsequently, RBF-based methods have been developed which address this issue. Carr et al. introduced a reconstruction method which requires the addition of "auxiliary" points to the data in an ε -width narrow band around the surface determined by (possibly approximated) normal vectors to the surface [9]. This method can be sensitive to the selection of the ε parameter, introducing numerical instabilities into the reconstruction, especially around thin features [33], and optimal choices for ε are lacking. While direct computation of this method can be expensive and requires solving a 3N-by-3N linear system, using fast summation algorithms [9], partition of unity [19], and compactly supported RBFs [37] have been shown to overcome this issue. The use of compactly supported RBFs has especially gained much attention, due to the resulting sparse linear systems; however, one must still choose a support radius for the compactly supported kernels, and if this value is too small, then the approximation power of the method can be impacted [33, 35, 37, 49]. CFPU addresses the issue of computational complexity with RBF systems while still remaining numerically stable and accurate. Additionally,

the method does not require the choice of shape parameter or support radii for the RBFs.

- 1.2. Contributions. In this paper, we present a novel method for reconstructing curves and surfaces with curl-free, vector-valued RBFs. CFPU is fast and requires only points on the surface and their corresponding normals. The RBFs we use are free of shape and support parameters, and the resulting linear systems are well-conditioned. Additionally, our method can be made interpolatory or can be regularized to handle noisy data. Since the implementation of the method involves local partition of unity patches, it is also scalable and highly parallelizable.
- 2. Curl-free RBFs. Curl-free RBFs were developed for the interpolation of curl-free vector fields that are given from scattered measurements as occurs, for example, in the areas of electrostatics and geodesy [1,16,21]. This technique has the important features that the vector interpolants are analytically curl-free and are well-posed for scattered data. Additionally, a scalar potential can be extracted directly from the interpolants that approximates the underlying potential of the field (up to an additive constant) [22]. Curl-free RBF interpolation is similar to scalar RBF interpolation in the sense that one constructs the interpolants from linear combinations of shifts of a kernel at each of the given data sites. The difference between the approaches is that in the curl-free case one uses shifts of a matrix-valued kernel whose columns are curl-free. For the sake of brevity, we do not review scalar RBF approximations but refer the reader to any of the books [19, 20, 51].

Let $\phi : \mathbb{R}^d \times \mathbb{R}^d \longrightarrow \mathbb{R}$ be a radial kernel in the sense that $\phi(\mathbf{x}, \mathbf{y}) = \eta(\|\mathbf{x} - \mathbf{y}\|)$, for some $\eta : [0, \infty) \longrightarrow \mathbb{R}$, where $\| \cdot \|$ is the vector two-norm, and $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. It is common in this case to simply write $\phi(\mathbf{x}, \mathbf{y}) = \phi(\|\mathbf{x} - \mathbf{y}\|)$ and refer to ϕ as an RBF. A matrix-valued curl-free kernel Φ is given as [21]

(2.1)
$$\Phi(\mathbf{x}, \mathbf{y}) = -\nabla \nabla^T \phi(\|\mathbf{x} - \mathbf{y}\|),$$

where ∇ is the gradient in \mathbb{R}^d applied to \mathbf{x} , and ϕ is assumed to have two continuous derivatives. Since Φ is built from an RBF, these kernels are simply called curl-free RBFs. For any $\mathbf{c} \in \mathbb{R}^d$ and fixed \mathbf{y} , the vector field $\Phi(\mathbf{x}, \mathbf{y})\mathbf{c}$ is curl-free in \mathbf{x} . This follows since

(2.2)
$$\Phi(\mathbf{x}, \mathbf{y})\mathbf{c} = \nabla \underbrace{\left(-\nabla^T \phi \left(\|\mathbf{x} - \mathbf{y}\|\right) \mathbf{c}\right)}_{g(\mathbf{x})} = \nabla (g(\mathbf{x})),$$

i.e. $\Phi(\mathbf{x}, \mathbf{y})\mathbf{c}$ is the gradient of a scalar function g. Note that the second argument of Φ acts as a shift of the kernel and indicates where the field (2.2) is "centered."

An interpolant to a curl-free vector field $\mathbf{u} \in \mathbb{R}^d$ sampled at distinct points $X = \{\mathbf{x}_j\}_{j=1}^N$ can be constructed from Φ as follows:

(2.3)
$$\mathbf{s}(\mathbf{x}) = \sum_{j=1}^{N} \Phi(\mathbf{x}, \mathbf{x}_j) \mathbf{c}_j,$$

where the expansion coefficients $\mathbf{c}_j \in \mathbb{R}^d$ are found by enforcing the interpolation conditions $\mathbf{s}|_X = \mathbf{u}|_X$. This results in the linear system

(2.4)
$$\sum_{i=1}^{N} \Phi(\mathbf{x}_i, \mathbf{x}_j) \mathbf{c}_j = \mathbf{u}_i, \ i = 1, 2, \dots, N,$$

which is commonly written as $A\mathbf{c} = \mathbf{u}$, where A is the block dN-by-dN interpolation matrix

(2.5)
$$A = \begin{bmatrix} \Phi(\mathbf{x}_1, \mathbf{x}_1) & \Phi(\mathbf{x}_1, \mathbf{x}_2) & \cdots & \Phi(\mathbf{x}_1, \mathbf{x}_N) \\ \Phi(\mathbf{x}_2, \mathbf{x}_1) & \Phi(\mathbf{x}_2, \mathbf{x}_2) & \cdots & \Phi(\mathbf{x}_2, \mathbf{x}_N) \\ \vdots & \vdots & \ddots & \vdots \\ \Phi(\mathbf{x}_N, \mathbf{x}_1) & \Phi(\mathbf{x}_N, \mathbf{x}_2) & \cdots & \Phi(\mathbf{x}_N, \mathbf{x}_N) \end{bmatrix}.$$

One can show that A is positive definite if Φ is constructed from an appropriately chosen scalarvalued ϕ [21]; see Table 1 for some examples. An important feature from the construction of the

Radial kernel	Expression	
Gaussian (GA)	$\phi(r) = \exp(-(\varepsilon r)^2)$	
Inverse multiquadric (IMQ)	$\phi(r) = (1 + (\varepsilon r)^2)^{-\frac{1}{2}}$	
Multiquadric (MQ)	$\phi(r) = -(1 + (\varepsilon r)^2)^{\frac{1}{2}}$	

Table

Examples of radial kernels that result in positive definite matrices A (2.5) for curl-free RBF interpolation. Here $\varepsilon > 0$ is the shape parameter.

curl-free RBF interpolant is that we can extract a scalar potential for the interpolated field by exploiting (2.2):

(2.6)
$$\mathbf{s}(\mathbf{x}) = \nabla \left(\underbrace{-\sum_{j=1}^{N} \nabla^{T} \phi \left(\|\mathbf{x} - \mathbf{x}_{j}\| \right) \mathbf{c}_{j}}_{s(\mathbf{x})} \right).$$

While the method described above will ensure a curl-free interpolant of the sampled field, some issues do arise. First, the size of the linear system (2.4) grows rapidly with N, and, for a globally supported kernel, will be dense and computationally expensive to solve—requiring $\mathcal{O}((dN)^3)$ operations if a direct method is used. Second, each evaluation of the interpolant (or potential (2.6)) involves dN terms, which can become computationally expensive when many evaluations are necessary (such as occurs in the present application). Other issues involve the shape parameter ε used in the radial kernels from Table 1. This parameter controls how flat or peaked the radial kernels are and has a dramatic effect on both the accuracy of the interpolant as well as the conditioning of the interpolation matrix A. If ε is fixed and the total number of interpolation points N grows, then the A matrix becomes exponentially ill-conditioned with N. Additionally, while extensive literature dedicated to finding the "good" values of ε to use exists for scalar RBF interpolation [8, 40], these approaches are computationally expensive, and this will only be exacerbated by the larger sizes of the linear systems for curl-free RBFs. To bypass these issues, we next discuss curl-free RBFs that do not feature a shape parameter. We address the issues with the computational cost in Section 3.

2.1. Curl-free polyharmonic splines. Interpolants based on polyharmonic splines (PHS) were introduced by Duchon as a generalization of univariate spline interpolants to higher dimensions [18]. PHS interpolants have the property that they minimize an energy functional that can be interpreted as a type of "bending energy" for the surfaces they produce, similar to univariate splines [51, ch. 13]. PHS are radial kernels and come in the following two types:

(2.7)
$$\phi_{\ell}(r) = (-1)^{\ell+1} \begin{cases} r^{2\ell} \log r, & \ell \text{ positive integer,} \\ r^{2\ell+1}, & \ell \text{ non-negative integer.} \end{cases}$$

For interpolation in \mathbb{R}^d , the first option is typically used for d even and the second option for d odd. The choice for the order parameter ℓ is often made based on smoothness assumptions of the data, with larger ℓ for smoother data. However, larger ℓ also negatively effects the numerical stability of the interpolants [51, ch. 12]. The combination of d and ℓ determines the minimization properties of the interpolants [51, ch. 13]; the choice of $\ell = 1$ for d = 2 leads to the classical thin-plate spline. PHS do not feature a shape parameter like other RBFs, as any scaling of r just factors out of the kernels. While ℓ is a free parameter, one does not need to continually search for a good value to use when the interpolation problem is changed, as is typically the case for RBFs with shape parameters.

Curl-free PHS were introduced in [1] and have further been studied in [4, 15]. These matrix-valued kernels can be produced by using (2.1) with ϕ_{ℓ} given by either of the choices in (2.7) and ℓ chosen large enough to ensure the derivatives make sense. We will denote these kernels by Φ_{ℓ} . As with scalar PHS, it is necessary to modify the curl-free RBF interpolant (2.3) to ensure a well-posed problem. In the curl-free PHS case, the interpolant (2.3) must be augmented with curl-free (vector) polynomials in \mathbb{R}^d of degree $\ell - 1$ [15], where degree refers to the total degree of any of the components of the (vector) polynomial. A basis for curl-free polynomials up to degree $\ell - 1$ can be generated as follows. Let $\{p_0, \ldots, p_L\}$ be a monomial basis for scalar polynomials up to degree ℓ in \mathbb{R}^d , where $L = {\ell+d \choose d} - 1$ and $p_0 = 1$. Then a basis for curl-free polynomials up to degree $\ell - 1$, $\{\mathbf{p}_1, \ldots, \mathbf{p}_L\}$, is given by applying the gradient to each p_i , i.e., $\mathbf{p}_i = \nabla p_i$, $i = 1, \ldots, L$. Note that the gradient reduces the degree of the scalar polynomials by one, which is why we end up with curl-free polynomials of degree $\ell - 1$. As an example, we give a basis for degree 1 curl-free polynomials in \mathbb{R}^2 and \mathbb{R}^3 obtained from a (scaled) basis of monomials of degree $\ell = 2$:

Poly. basis in
$$\mathbb{R}^2$$
: $\left\{ \begin{bmatrix} 1\\0 \end{bmatrix}, \begin{bmatrix} 0\\1 \end{bmatrix}, \begin{bmatrix} y\\x \end{bmatrix}, \begin{bmatrix} x\\0 \end{bmatrix}, \begin{bmatrix} 0\\y \end{bmatrix} \right\}$
Poly. basis in \mathbb{R}^3 : $\left\{ \begin{bmatrix} 1\\0\\0 \end{bmatrix}, \begin{bmatrix} 0\\1\\0 \end{bmatrix}, \begin{bmatrix} 0\\0\\1 \end{bmatrix}, \begin{bmatrix} y\\x\\0 \end{bmatrix}, \begin{bmatrix} y\\x\\0 \end{bmatrix}, \begin{bmatrix} z\\0\\x \end{bmatrix}, \begin{bmatrix} 0\\z\\y \end{bmatrix}, \begin{bmatrix} x\\0\\0 \end{bmatrix}, \begin{bmatrix} 0\\y\\0 \end{bmatrix}, \begin{bmatrix} 0\\0\\z \end{bmatrix} \right\}$.

A curl-free PHS interpolant of order ℓ to a curl-free vector field $\mathbf{u} \in \mathbb{R}^d$ sampled at distinct points $X = \{\mathbf{x}_j\}_{j=1}^N$ is given as follows:

(2.8)
$$\mathbf{s}(\mathbf{x}) = \sum_{j=1}^{N} \Phi_{\ell}(\mathbf{x}, \mathbf{x}_{j}) \mathbf{c}_{j} + \sum_{k=1}^{L} b_{k} \mathbf{p}_{k}(\mathbf{x}),$$

where the interpolation coefficients $\mathbf{c}_j \in \mathbb{R}^d$ and $b_k \in \mathbb{R}$ are determined by the conditions $\mathbf{s}|_X = \mathbf{u}|_X$ as well as the constraints

$$\sum_{j=1}^{N} \mathbf{c}_{j}^{T} \mathbf{p}_{k}(\mathbf{x}_{j}) = 0, \qquad k = 1, 2, \dots, L.$$

These constraints are necessary for the interpolant to minimize a certain energy norm and also limit its far-field growth [15]. We can state both constraints in terms of the following linear system of equations:

(2.9)
$$\begin{bmatrix} A & P \\ P^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{c} \\ \mathbf{b} \end{bmatrix} = \begin{bmatrix} \mathbf{u} \\ \mathbf{0} \end{bmatrix},$$

where A is defined in (2.5) and

$$P = \begin{bmatrix} \mathbf{p}_1(\mathbf{x}_1) & \mathbf{p}_2(\mathbf{x}_1) & \cdots & \mathbf{p}_L(\mathbf{x}_1) \\ \mathbf{p}_1(\mathbf{x}_2) & \mathbf{p}_2(\mathbf{x}_2) & \cdots & \mathbf{p}_L(\mathbf{x}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{p}_1(\mathbf{x}_N) & \mathbf{p}_2(\mathbf{x}_N) & \cdots & \mathbf{p}_L(\mathbf{x}_N) \end{bmatrix}.$$

Provided the set of points X is unisolvent with respect to the curl-free polynomial basis (i.e. P is full rank), this linear system is non-singular and thus the interpolation problem is well-posed [15]. However, as with interpolation matrices based on curl-free RBFs with shape parameters, this interpolation matrix also becomes ill-conditioned as N increases, but the growth is algebraic as opposed to exponential [21].

We note that a scalar potential s can also be recovered similarly to (2.6) as

(2.10)
$$\mathbf{s}(\mathbf{x}) = \nabla \left(\underbrace{-\sum_{j=1}^{N} \nabla^{T} \phi_{\ell} (\|\mathbf{x} - \mathbf{x}_{j}\|) \mathbf{c}_{j} + \sum_{k=1}^{L} b_{k} p_{k} (\mathbf{x})}_{s(\mathbf{x})} \right),$$

where p_k are the scalar polynomial basis used to generate the curl-free polynomial basis. This potential plays a key role in the CFPU method.

2.2. An example. At this point, it is illustrative to see how curl-free RBFs can be used to recover a level set from an oriented point cloud. We focus here on the case of a level curve $\mathcal{P} \subset \mathbb{R}^2$, since this is the situation in which the global method described above would be applicable due to the smaller problem sizes. The example we focus on uses the Cassini oval as the target level curve to recover, which can be described as the zero-level set of the implicit function

(2.11)
$$f(\mathbf{x}) = f(x_1, x_2) = (x_1^2 + x_2^2)^2 - 2a^2(x_1^2 - x_2^2) + a^4 - b^4,$$

where we take a=1 and b=1.1; see the Figure 1(b) for the resulting curve (dashed-line). We sample this curve at N nodes $X = \{\mathbf{x}_j\}_{j=1}^N$ and compute the corresponding (unit) normal vectors using $\mathbf{n}_j = \nabla f(\mathbf{x}_j)$; see Figure 1(a) for a plot of the exact data used for the case of N=30. As

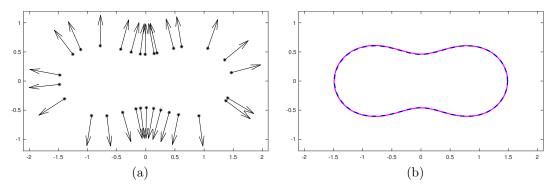


FIGURE 1. (a) N=30 points sampled from a Cassini oval (2.11) with a=1 and b=1.1, together with the corresponding normal vectors to the curve. (b) Zero-level curve from the global curl-free PHS interpolation method with $\ell=2$ (magenta) together with the exact curve \mathcal{P} (blue dashed line).

mentioned in the introduction, the key to our method is the fact that the normal vectors to a level curve (or surface) are curl-free. We thus fit the data $(\mathbf{x}_j, \mathbf{n}_j)$, $j = 1, \ldots, N$, using the curl-free RBF interpolant (2.8) and from this extract out the potential s as in (2.10). Since the potential for a curl-free field is only unique up to a constant (which is a consequence of the Helmholtz decomposition theorem [22]), the zero-level curve of s will not necessarily approximate the zero-level curve of s this we set $\tilde{s}(\mathbf{x}) = s(\mathbf{x}) - \mu$, where μ is the discrete mean of s at the nodes s. The result from this experiment is shown in Figure 1(b), where we see excellent agreement between the zero-level curve of s and s.

While this global method is reasonable to use for reconstructing level curves and surfaces when N is small, the cost of solving the linear system (2.9) becomes too expensive for large N, which will be the case for any sampling of a complex surface. We address this issue next.

3. The CFPU method. Partition of unity (PU) methods offer a way to split up a global approximation problem on a domain Ω into local approximation problems on overlapping patches covering Ω . These local approximations are then blended together to form a global approximant

using weight functions that form a partition of unity [2]. This procedure can drastically reduce the computational cost of the original approximation problem. In order to explain the CFPU method, we first give a brief description of the ideas behind PU methods as they pertain to our problem, and introduce some necessary notation for what follows.

3.1. PU methods. Let $\Omega \subset \mathbb{R}^d$ be an open, bounded domain of interest for approximating some function $f:\Omega \longrightarrow \mathbb{R}$. Let Ω_1,\ldots,Ω_M be a collection of distinct overlapping patches that form an open cover of Ω , i.e., $\bigcup_{m=1}^M \Omega_m \supseteq \Omega$, and let the overlap between patches be limited such that at most $K \ll M$ patches overlap at any given point $\mathbf{x} \in \Omega$. For each $m=1,\ldots,M$, let $w_m:\Omega_m \longrightarrow [0,1]$ be a weight function such that w_m is compactly supported on Ω_m , and let the set of weight functions $\{w_m\}$ have the property that $\sum_{m=1}^M w_m \equiv 1$. Suppose s_m is some approximation to f on each patch Ω_m . Then the PU approach of Babuška and Melenk [2] is to form an approximant s to f over the whole domain Ω by "blending" the local approximants s_m with w_m as follows: $s = \sum_{m=1}^M w_m s_m$.

Balls are common choices for the patches, since then the weight functions w_m can be easily constructed using Shepard's method [43] as follows. Let $\kappa : \mathbb{R}^+ \to \mathbb{R}$ have compact support over the interval [0, 1). For each patch Ω_m , let $\boldsymbol{\xi}_m \in \mathbb{R}^d$ denote its center and $\rho_m > 0$ denote its radius, and define $\kappa_m(\mathbf{x}) := \kappa (\|\mathbf{x} - \boldsymbol{\xi}_m\|/\rho_m)$. The weight function for patch Ω_m is then given by

(3.1)
$$w_m(\mathbf{x}) = \kappa_m(\mathbf{x}) / \sum_{j=1}^M \kappa_j(\mathbf{x}), \ m = 1, \dots, M.$$

Note that each w_m is only supported over Ω_m and that the summation in the denominator only involves terms that are non-zero over patch Ω_m , which is bounded by K. In this study, we choose the patches as balls, and use the C^1 quadratic B-spline

(3.2)
$$\kappa(r) = \begin{cases} 1 - 3r^2, & 0 \le r \le \frac{1}{3}, \\ \frac{3}{2}(1 - r)^2, & \frac{1}{3} \le r \le 1 \end{cases}$$

to define the weight functions. With this weight function the global approximant s is at most C^1 . However, weight functions with more smoothness can be easily created from smoother compactly supported univariate functions κ if global approximants with higher smoothness are desired.

RBFs are commonly used with the PU approach to reduce the computational cost. They have been used for approximating a function from scattered samples (e.g. [11, 12, 50]) and solving differential equations (e.g. [31, 41, 42]). Recently the present authors presented a PU method for interpolation of divergence-free and curl-free vector fields [17], which is what the current approach is based on. For RBF-PU methods, it should be noted that the smoothness of the weight functions (3.1) does not govern the rates of convergence of the global PU approximants [17, 31].

3.2. Description of CFPU. For brevity, we describe the CFPU method for reconstructing a zero-level surface \mathcal{P} in \mathbb{R}^3 defined by $f(\mathbf{x}) = 0$ using curl-free PHS of order ℓ . Let $X = \{\mathbf{x}_j\}_{j=1}^N$ be a given set points on \mathcal{P} and let $\{\mathbf{n}_j\}_{j=1}^N$ denote the unit normals (or approximations to the normals) of \mathcal{P} at X. Let $\Omega_1, \ldots, \Omega_M$ be a set of overlapping patches that form an open cover of \mathcal{P} . Finally, let X_m denote the nodes contained in Ω_m and n_m denote the cardinality of X_m . For each Ω_m , we fit a curl-free RBF interpolant \mathbf{s}_m of the form (2.8) to the normals at the points in X_m and then extract from this its scalar potential s_m using (2.10).

A natural first approach to constructing a global potential for approximating \mathcal{P} would be to blend these local potentials s_m using the PU weight functions w_m as $s = \sum_{m=1}^M w_m s_m$. However, this will lead to an issue since each s_m is only unique up to a constant. This means that for two patches Ω_k and Ω_m that overlap, s_k and s_m may be shifted from one another in the overlap region,

which would then lead to an inaccurate global PU approximant in the overlap. To fix this issue we can shift each potential by a different constant so that they approximately agree in the overlap region. To determine these constants, we use the fact that the points X_m reside on the zero-level surface \mathcal{P} , and we want each potential s_m to be approximately zero on X_m . One way to achieve this result is to enforce that the discrete mean of the local potentials is zero over the patch nodes. To this end, let μ_m equal the mean of the values of s_m at X_m , and then define the shifted potential

$$\overline{s}_m := s_m - \mu_m$$
.

The global CFPU approximant for the underlying implicit function f is given by

(3.3)
$$s(\mathbf{x}) := \sum_{m=1}^{M} w_m(\mathbf{x}) \overline{s}_m(\mathbf{x}).$$

We can then approximate \mathcal{P} as the surface defined by the set of all \mathbf{x} in $\bigcup_{m=1}^{M} \Omega_m$ such that $s(\mathbf{x}) = 0$. The CFPU method requires solving M linear systems of size $(3n_m + L)$ -by- $(3n_m + L)$ rather than one large (3N + L)-by-(3N + L) system for the global curl-free method described in Section 2.2. We select $n_m \ll N$, for all m so that the computational cost is significantly reduced, i.e. the complexity is $\mathcal{O}(\sum_{m=1}^{M} (3n_m + L)^3 N)$ rather than $\mathcal{O}((3N + L)^3)$, when using a direct solver. Furthermore, each of these smaller systems can be solved independently, making the CFPU method pleasingly parallel compared to the global method. Finally we note that the computational complexity for evaluating the CFPU approximant (3.3) is also significantly less than the global method. For each evaluation point, only a small subset (equal to the number of patches that contain the evaluation point) of the local potentials \bar{s}_m need to be evaluated. The cost of evaluating each \bar{s}_m is $\mathcal{O}(3n_m + L)$ rather than $\mathcal{O}(3N + L)$ for the global method. These potentials can also be evaluated independently.

3.3. Exact interpolation. The CFPU method as described above will not in general exactly interpolate the zero-level surface \mathcal{P} at the points in X, i.e. $s(\mathbf{x}_j) \neq 0$, j = 1, ..., N, which is often desirable when the points are assumed to be exactly on the surface [34]. We can, however, enforce this condition by simply subtracting an interpolant of the residual from each patch potential s_m . We describe the details of this procedure below.

Let the points in patch Ω_m from X be denoted by $X_m = \{\mathbf{x}_j^m\}_{j=1}^{n_m}$ and let σ_m be a scalar PHS interpolant to the values of s_m at X_m . Using the same notation from Section 2.1, this interpolant can be written as

(3.4)
$$\sigma_m(\mathbf{x}) = \sum_{j=1}^{n_m} c_j^m \phi_\ell(\|\mathbf{x} - \mathbf{x}_j^m\|) + \sum_{k=0}^L b_k^m p_k(\mathbf{x}),$$

where the coefficients are determined from the interpolation conditions $\sigma_m|_{X_m} = s_m|_{X_m}$ and the moment conditions $\sum_{j=1}^{n_m} c_j^m p_k(\mathbf{x}_j^m) = 0$, $k = 0, \ldots, L$. The potential s_m on patch Ω_m can be shifted by σ_m to obtain

$$\widetilde{s}_m := s_m - \sigma_m.$$

By construction, this shifted potential satisfies $\tilde{s}_m(\mathbf{x}_j^m) = 0$, so that an interpolatory global CFPU approximant for the underlying implicit function f can then be obtained as

(3.5)
$$s(\mathbf{x}) := \sum_{m=1}^{M} w_m(\mathbf{x}) \widetilde{s}_m(\mathbf{x}).$$

An approximation to the surface \mathcal{P} is again given as the set of all \mathbf{x} in $\bigcup_{m=1}^{M} \Omega_m$ such that $s(\mathbf{x}) = 0$.

We note that this exact interpolation technique is more expensive than just shifting the potentials by the mean as in (3.3), but only by a constant factor. Additionally, just as with s_m , each σ_m can be determined independently of the others.

The residual of the patch potentials s_m can be highly oscillatory, which could lead to spurious oscillations in the interpolants σ_m and hence also s. For 3D reconstructions, we thus recommend using the PHS kernel ϕ_0 (i.e. $\phi_0(r) = r$) for the residual since the function will have minimal bending energy among all interpolants [18]. This is the choice we use in all of our examples.

- **3.4.** Regularization. There are two regularization techniques we can apply to deal with noise.
- **3.4.1.** Normals. If the samples of the normal vectors of \mathcal{P} are corrupted with noise, then interpolating them exactly on each patch to recover the potentials may cause issues, such as producing spurious sheets in the reconstructed surface. In this case, it may make sense to instead introduce some regularization in the vector approximants on the patches. Regularized kernel approximation, such as smoothing splines or ridge regression [48], offers one effective way to do this.

For a curl-free PHS approximant given in (2.8), the smoothing spline regularization approach amounts to solving the following minimization problem:

(3.6)
$$\min_{\mathbf{c} \in \mathbb{R}^{3n}} \left[\frac{1}{3n} \sum_{j=1}^{n} \|\mathbf{s}(\mathbf{x}_j) - \mathbf{n}_j\|^2 + \lambda \mathbf{c}^T A \mathbf{c} \right], \text{ subject to } P^T \mathbf{c} = 0,$$

where A and P are the matrices from (2.9) for n points. The first term in the quadratic functional measures the goodness of fit of the approximant while the second term measures its smoothness.¹ The regularization parameter $\lambda \geq 0$ controls the tradeoff between these terms, with larger λ resulting in smoother approximants. For a given λ , we can obtain the minimizer of the constrained quadratic functional (3.6) by solving the following modified version of the system (2.9):

(3.7)
$$\begin{bmatrix} A + 3n\lambda I & P \\ P^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{c} \\ \mathbf{b} \end{bmatrix} = \begin{bmatrix} \mathbf{u} \\ \mathbf{0} \end{bmatrix},$$

where I is the 3n-by-3n identity matrix and \mathbf{u} contains the normals.

For the CFPU method, we use this regularization approach on each patch Ω_m to obtain regularized potentials s_m . This opens up the option of using a different regularization parameter λ_m on each patch, and thus controlling the regularization of the approximants spatially.

In Section 4, we present some results for selecting the regularization parameters λ_m using generalized cross validation (GCV) [48] and use the method from [3] to compute these parameters. The GCV approach is more computationally expensive than an ad hoc approach of choosing λ_m and then solving (3.7) directly, but it does provide a means for automating the choice of parameters.

3.4.2. Residual. If the point samples in the point cloud are also noisy, as often occurs in range scans of real 3D objects, then enforcing exact interpolation on the patch potentials by interpolating the residual may again cause issues in the reconstruction. We can also introduce regularization in this process using smoothing splines. In this case, one uses a similar minimization problem as (3.6), but for the scalar TPS approximant (3.4). In fact, smoothing splines were developed for this scalar problem [48]. Since different regularization parameters can be used for fitting the normals vs. fitting the residual of the potential, we let α denote the residual regularization parameter to avoid confusion. We also experiment with picking this parameter using GCV in Section 4.

¹This term arises from the minimization of a Hilbert space semi-norm related to the function space of the vector approximants—the so called native space norm associated with the Φ kernel [4].

3.5. Choosing the PU patches. As discussed above, each PU patch Ω_m , $m=1,\ldots,M$, is chosen as a ball and can thus be determined by its center $\boldsymbol{\xi}_m$ and radius ρ_m . In this work, we choose the set of patch centers $\Xi = \{\boldsymbol{\xi}_m\}_{m=1}^M$ directly from the point cloud using the point cloud simplification technique implemented in MeshLab [13], which is based on Poisson disk sampling [14]. The input to the function is an estimate for the number of subsamples of the point cloud that should be produced. The result is a quasi-uniform point cloud with respect to the (approximated) geodesic distance for the surface represented by the original point cloud. We use this reduced point cloud for the patch centers. The point cloud simplification technique is computationally efficient, with an estimated complexity of $\mathcal{O}(M \log M)$. Another option for generating the patch centers is to use the weighted sample elimination method from [53], which is also a Poisson disk sampling technique. This method allows one to produce an exact number for the reduced point cloud, but is not as computationally efficient as the method in MeshLab.

The radii for the patches are determine using the following procedure. For each patch center $\boldsymbol{\xi}_m \in \Xi$, we determine the distance to its nearest neighbor (using the Euclidean distance) in the set $\Xi \setminus \boldsymbol{\xi}_m$ using a k-d tree. We then set τ as the maximum of these distances. An initial radius ρ for all the patches is then chosen as $\rho = (1+\delta)\tau/2$, where $\delta > 0$. The parameter δ controls the overlap of the patches and we have found that using $\delta = 1$ gives consistently good results. Setting $\rho_m = \rho$ initially for each patch Ω_m , we then check if the number of samples, n_m , it contains satisfies $n_m \geq n_{\min}$. If it does not, then we increase ρ_m until this inequality is satisfied. In this work, we choose $n_{\min} = 2L = 2(\binom{\ell+3}{3}-1)$, where ℓ is the order of the curl-free PHS kernel used. It is necessary for the curl-free PHS interpolant to use L points to be well-posed, and we have found that doubling this value works well in practice. The final step involves adjusting the radii of the patches so that every sample point is included in at least one patch. We do this by finding the closest patch center to each outlying point and increasing its radius so that point is enclosed.

In our implementation, we build a k-d tree on the patch centers and use this for range queries in determining which sample points belong to which patches.

4. Results. In this section, we test the CFPU method on several different 3D surface reconstruction problems. We start with a comparison of the CFPU method with the global curl-free method for a small point cloud. We then focus solely on the CFPU method. We use point clouds generated from a known surface to test the accuracy of the method. Using this same known surface, we then add noise to the normals and show how regularization can help reconstruct smooth surfaces with noisy normals. Next, we consider a problem with raw range data that contains misalignments of the points and noise to show how both regularization of the normals and the residuals can result in a smooth reconstruction. We then show how the method performs for reconstructing various common surfaces found in the literature. Lastly, we present results on computational performance of the CFPU method. For many of the examples, we also include comparisons to the surface reconstructions from two popular methods, the Screened Poisson Surface Reconstruction (SPR) method [26,27] and the Smooth Signed Distance (SSD) method [7]. Unlike many other methods (e.g. HRBF Implicits), open-source versions of SPR and SSD are freely available [28]. Table 2 lists the names of the surfaces we use in the experiments together with the size of the point cloud and the number patches M used for the CFPU reconstructions. The choices for M were chosen in a somewhat ad hoc manner and were not optimized.

While the problems we consider have normals that can be either computed analytically or that are supplied with the point cloud data, this may generally not be the case. There are many algorithms available that can approximate the normals directly from a point cloud data (e.g., [25, 30, 39]). In some of the results that follow, we have experimented with the point cloud normal estimation function from MeshLab. The last column of Table 2 lists the surfaces for which we used this function.

In all of the tests, we use either exact interpolation or its regularized version to approximate the potentials s (cf. (3.5)). We found that both of these procedures result in better reconstructions

Surface	Num. points N Num. patches M		Normals
Homer	5103	135	MeshLab
Trefoil	[6114, 32856]	[6114,32856] 864	
Stanford Bunny	362271	14472	Raw
Stanford Dragon	434856	14400	Raw
Happy Buddha	583079	42861	MeshLab
Armadillo	172974	14349	MeshLab
Raptor	135740	10337	Raw
Filigree	514300	35130	Raw
Pump Carter	238141	14351	Raw
Dancing children	724588	17296	Raw
Gargoyle	499548	21436	Raw

Table 2

Various properties and parameters used in the numerical examples. Entries in the Normals column indicate the technique for generating the normals: Exact=computed exactly from the surface; MeshLab=computed using the "Compute normals for point set" in MeshLab using 10 neighbors and no smoothing; Raw=normals are used directly from the values given in the downloaded files from either the Stanford University 3D Scanning Repository or the AIM@SHAPE-VISIONAIR repository.

of the zero-level surface than the mean-shift version (cf. (3.3)). Additionally, we use the curl-free PHS kernels generated from the $\phi_{\ell}(r) = (-1)^{\ell+1} r^{2\ell+1}$ since the point clouds are in \mathbb{R}^3 . As discussed previously, the order ℓ controls the smoothness of the kernels.

All results were obtained from a MATLAB implementation of the CFPU method. Additionally, we used the MATLAB isosurface function to extract the zero-level set for the surfaces based on a sufficiently dense sampling on a uniform 3D grid surrounding the point cloud. The code is freely available from the CFPU GitHub repository [52], including MEX interfaces to the MeshLab functions (VCGLib library) used determining the patch centers and normals. All the CFPU results presented below can be reproduced from the code and data in the repository.

4.1. Comparison with the global method. We first compare the CFPU method with the global curl-free method discussed in Section 2.2, but now for reconstructing a 3D surface. As mentioned previously, the global method is prohibitively expensive for large data sets, so we restrict our example to the small Homer point cloud shown on the left of Figure 2, which consists of N=5103 points. The resulting reconstructions using $\ell=1$ and 2 for the two curl-free methods are shown in the second and third columns of Figure 2, respectively. We can see from the figure that for $\ell=1$, both methods produce visually pleasing, essentially indistinguishable level-surfaces from the point cloud. However, when using the smoother $\ell=2$ kernel, both methods result in spurious sheets. While these can be resolved by including some regularization of the normals, we do not present the results here, as later tests will experiment with this procedure.

Reconstructions of Homer using the SPR and SSD methods are displayed in the fourth column of Figure 2. We see from the figure that these methods result in rougher, less sharp surfaces than the curl-free methods with $\ell = 1$. This may be due to the point cloud having sparse sampling in certain regions [27]. The curl-free methods seem to be less sensitive to these issues.

4.2. Accuracy of CFPU. We consider a pipe surface of radius 0.7 generated from a (2,5) torus knot. The 3D curve defining the knot can be written parametrically as

$$(x(t), y(t), z(t)) = (\cos(2t)(\cos(5t) + 3), \sin(2t)(\cos(5t) + 3), \sin(5t)),$$

where $0 \le t \le 2\pi$. From this parametric representation, we generate a point cloud $X = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ and the (unit) normals for the knot (see Figure 3 for an illustration), and then use these in the CFPU

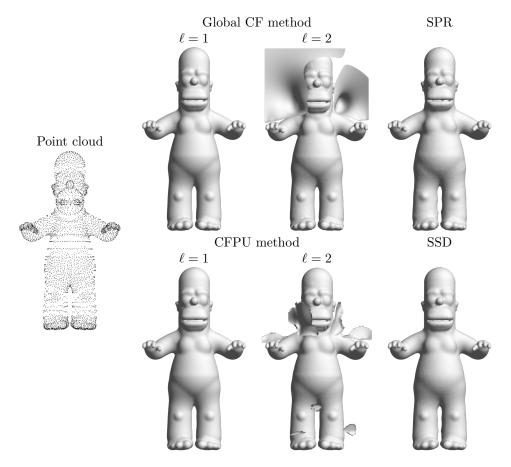


Figure 2. Comparisons of different reconstructions of Homer.

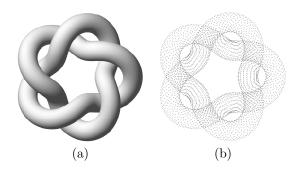


FIGURE 3. (a) Illustration of the knot used in the accuracy and noisy normals tests. (b) Example of a synthetic point cloud for the knot with N=6144 samples used in the tests.

method to reconstruct a potential whose zero-level surface approximates the exact knot surface. To test the accuracy of these reconstructions, we generate a dense set of 131424 points exactly on the knot and compute the root mean square (RMS) of the differences between the potential evaluated at these points and the exact solution (which is zero).

Table 3 displays the results for this test for $\ell=1$ and $\ell=2$ and for increasing numbers of sample points. In these tests, the synthetic point clouds were generated so that the average spacing

	$\ell = 1$		$\ell = 2$		
N	RMS error	Rate	RMS error	Rate	
6144	2.92×10^{-4}	_	1.88×10^{-5}	_	
8664	1.67×10^{-4}	3.26	8.60×10^{-6}	4.56	
11616	1.09×10^{-4}	2.88	4.21×10^{-6}	4.87	
18816	5.05×10^{-5}	3.20	1.23×10^{-6}	5.12	
23064	3.80×10^{-5}	2.81	7.46×10^{-7}	4.89	
27744	2.88×10^{-5}	2.99	4.73×10^{-7}	4.92	
32856	2.19×10^{-5}	3.25	3.08×10^{-7}	5.07	

Table 3

Comparison of the errors and estimated algebraic convergence rates in the CFPU reconstruction of the knot for increasing numbers of samples N using the curl-free PHS kernel Φ_{ℓ} , for $\ell=1,2$. All results use a fixed number of M=864 PU patches.

between points decreases like $N^{-1/2}$. So, we also include estimates of the algebraic convergence rates ν of the reconstructions assuming the RMS error is $\mathcal{O}(N^{-\nu/2})$. We see from the table that for both $\ell=1$ and $\ell=2$, the CFPU reconstructions appear to be converging to the true zero-level surfaces as the density of the samples increases. Furthermore, the reconstructions using $\ell=2$ appear to be converging at an approximate rate of $\nu=5$, while they appear to be converging at the approximate rate of $\nu=3$ for $\ell=1$. The faster rate is expected for $\ell=2$ since the knot is a smooth surface and higher values of ℓ produce smoother approximants.

More detailed results on the accuracy in approximating general curl-free fields (and their potentials) in two and three dimensions with curl-free RBF interpolants can be found in [17]. For example, it is shown that for certain piecewise smooth radial kernels, the approximation error typically decreases when more points are enclosed in each patch. However, the rates of convergence remain the same.

4.3. Noisy normals. In this test we demonstrate how the regularization method from Section 3.4.1 can help with noise in the normals; we do not test the residual regularization method here. We use the knot example from the previous section and add noise to the exact (unit) normals $\{\mathbf{n}_1,\ldots,\mathbf{n}_N\}$ according to $\mathbf{n}_j^* = \mathbf{n}_j + \boldsymbol{\epsilon}_j$, where $\boldsymbol{\epsilon}_j \in \mathbb{R}^3$ and each of its components is an independent normally distributed random variable with mean zero and standard deviation 0.3. The results are displayed in Figure 4.

The first column of this figure shows the reconstructed surfaces for both $\ell=1$ and $\ell=2$ when no regularization is applied. Both results give a noisy surface, but only the $\ell=2$ results give spurious sheets. The second and third columns show the results when a fixed regularization parameter λ is used for all patches. We see from the figure that including regularization results in a smoother surfaces and increasing λ increases the smoothness. These columns also show that larger values of λ are needed for $\ell=1$ to obtain the same level of smoothness. The fourth column of the figure shows the results when automating the selection of the regularization parameters using GCV on each patch. We see from the figure that this technique appears to be quite effective at reducing the noise for both $\ell=1$ and $\ell=2$. However, in the $\ell=1$ case, there are still some remnants of oscillations where the knot crosses over itself.

For comparison purposes, the last column of Figure 4 shows reconstructions of the "noisy" knot using the SPR and SSD methods. To filter out the noise and produce these surfaces, we had to modify the default parameters to introduce more smoothing, similar to the CFPU method. For the SPR method we set the pointWeight parameter to 10 and the samplesPerNode parameter to 2, while for the SSD method we set the samplesPerNode parameter to 4 and the biLapWeight parameter to 4. We found these parameters produced the best results. We can see from the figure

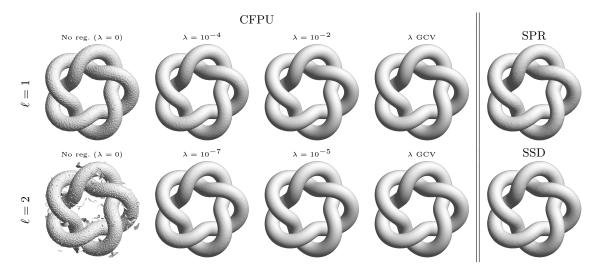


FIGURE 4. Comparison of reconstructions of the knot with zero mean Gaussian white noise added to the normals. First column shows the CFPU reconstructions without any regularization of the normals. Second and third columns show the CFPU reconstructions using regularization with a fixed parameter λ chosen for all the patches. Fourth column shows the CFPU reconstructions with the regularization parameter chosen using GCV on each patch. Fifth column shows the surfaces from the SPR and SSD methods for comparison purposes. All results use N=23064.

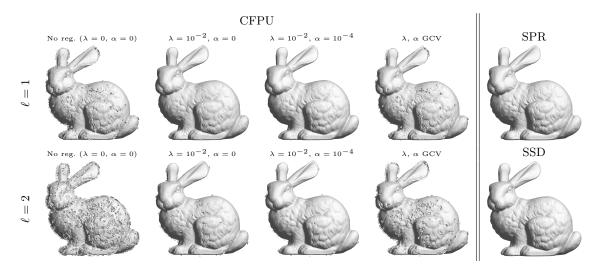


FIGURE 5. Comparison of the CFPU reconstructions of the Stanford Bunny using raw range data. First column shows the reconstructions without any regularization. The remaining columns show the effects of regularizing the fit of the normal vectors $(\lambda > 0)$ and the fit of the residual of the potential $(\alpha > 0)$. For the second and third columns, the parameters were fixed for all patches, while for the fourth column, the parameters were chosen on a per-patch basis using GCV. The last column shows the surfaces from the SPR and SSD methods for comparison purposes.

that both the SPR and SSD reconstructions produce more noisy surfaces than the regularized CFPU method, with the noise being more prevalent in the former.

4.4. Raw range data. We now consider reconstructing a surface from range data, which typically consists of optical scans of an object from different orientations and produces data sets that are misaligned and contain noise. For this test, we use the Stanford Bunny problem and

generate the point cloud and normals by combining the 10 available range scans. Figure 5 shows the resulting CFPU reconstructions of the surface for $\ell=1$ and $\ell=2$ and different regularization parameters.

We see from the first column of the figure that without any regularization, the CFPU method produces a zero-level surface with several spurious sheets and that these are more prevalent in the $\ell=2$ reconstruction than the $\ell=1$. The second column shows the reconstruction from regularizing the fit of the normal vectors using the same parameter for each patch. We see that this regularization reduces the spurious sheets for the $\ell=2$ case and completely eliminates them for the $\ell=1$ case. The third column shows the results of using both regularization of the normals and the fit of the residual of the potential, again using the same parameter for all the patches. For both values of ℓ , we see that this added regularization results in a smoother surface with less fine-scale detail. For the $\ell=2$ case, there are still spurious sheets present in the reconstruction and further tuning of the parameters could not eliminate these. Finally, the fourth column of Figure 5 shows the reconstructions when using GCV to determine both regularization parameters on a per-patch basis. Unlike the knot problem, we see from the figure that GCV is not effective at producing parameters that smooth the results and eliminate the spurious sheets. This may be due to the noise in the data being non-Gaussian and from inconsistent noise levels across the different scans.

The final column of Figure 5 shows the reconstructions of the bunny using the SPR and SSD using the default parameters. These surfaces are similar in terms of fine-scale detail and lack of spurious sheets to the CFPU surface with $\ell=1$ in the second column.

4.5. Common 3D models. We finally test the CFPU method for reconstructing some common 3D models found in the literature. We only present results for the method using $\ell=1$ since this gave the best reconstructions and always produced zero-level surfaces without spurious sheets. We also only include comparisons to the SPR method, since the SSD method gave similar results.

In the first set of experiments, we consider reconstructing three different CAD models, the Raptor, Filigree, and Pump Carter, using only the vertices and normals from the triangulations that define the models. In the case of the Pump Carter, the point cloud is obtained from a uniform subdivision of the original CAD model. The results of the CFPU reconstructions are displayed in the top row of Figure 6, while the bottom row displays the SPR reconstructions. No regularization was used for the CFPU results and the default parameters were used for the SPR results. We can see from the figure that CFPU and SPR produces very similar results. The fine-scale detail of these models is captured and there are no spurious sheets. Furthermore, the sharp features present in the models are well captured without over smoothing.

For the next set of experiments, we consider three models produced from scans of real objects, the Stanford Dragon, Happy Buddha, and Armadillo. In these tests, we use point clouds generated from post-processed versions of the scans where any outliers in the range data have been removed and the point clouds have been simplified. The top row of Figure 7 shows the resulting CFPU reconstructions of these models, while the bottom row shows the SPR reconstructions. As with the CAD models, no regularization was used for the CFPU results and the default parameters were used for the SPR results. We can see from the figure that the CFPU and SPR methods again produce similar results. The fine-scale detail is well-captured, without any noticeable smoothing, and there are no spurious sheets.

In the final set of experiments, we again consider post-processed models produced from scans of real objects. However, unlike the previous experiments, the CFPU reconstructions produce unwanted effects in some very localized areas of these models. Figure 8 displays the results for the Gargoyle and Dancing Children models. The first row of the figure shows the CFPU reconstructions without any regularization. The insets of these plots show a more detailed view of where some spurious sheets are produced in the reconstructions. These are exactly in the regions where the original polygonal models also have issues. The next row of the figure shows the results when including regularization

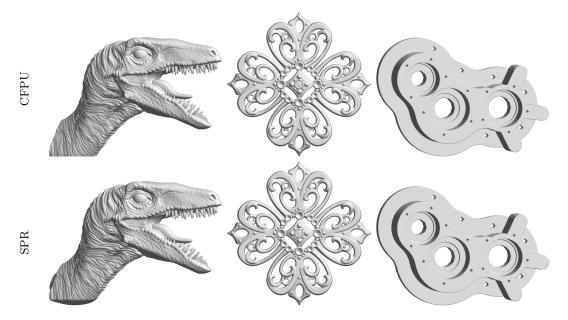


FIGURE 6. Surface reconstructions for the CFPU (top row) and SPR (bottom row) methods for three different surfaces produced from CAD models (Raptor, Filigree, and Pump Carter).

		Time in seconds			
Surface	Grid size	Number of cores			
		1	2	4	8
Raptor	$384 \times 143 \times 258$	18.2 (16.5)	6.67 (8.84)	4.14 (7.22)	3.35 (5.69)
Filigree	$383 \times 63 \times 384$	30.4 (7.70)	10.9 (4.10)	6.63(3.39)	4.99 (3.13)
Pump Carter	$384 \times 85 \times 271$	18.1 (12.4)	6.79(6.01)	4.22(4.68)	3.55 (4.11)
Stanford Dragon	$384 \times 175 \times 272$	42.73 (11.83)	17.24 (6.46)	10.7(5.46)	9.20 (4.82)
Happy Buddha	$161 \times 162 \times 384$	37.34 (7.71)	14.0 (4.20)	8.32 (3.43)	5.99 (3.02)
Armadillo	$323 \times 294 \times 384$	9.42 (10.0)	3.62(5.54)	2.25(4.85)	1.61 (4.55)

Table 4

Runtime performance of the CFPU method for reconstructing various surfaces. The first number in each of the timing columns is the time it takes to solve all the local linear systems for determining the expansion coefficients of every \widetilde{s}_m in (3.5). The number in parenthesis is the time it takes to evaluate the potential s in (3.5) over the grid points to extract the level surface. Refer to Table 2 for further details on the parameters used.

in the reconstructions across all the patches. For the Gargoyle model, we found that the spurious sheets around the front teeth could be eliminated by simply regularizing the fit of the normal vectors. However, for the Dancing Children model it was necessary include some regularization of both the normals and the residual to eliminate the spurious sheets near the joined hands. As such, we see that more smoothing has been introduced in Dancing Children results from the non-regularized version.

Since the issues in these models is localized to a very small area, it is possible to only use regularization in the patches that surround these areas. The last row of Figure 8 shows the results of such a localized regularization strategy. We can see from these plots that the spurious sheets have been eliminated with this approach while not over-smoothing the entire model. This is especially prevalent in the Dancing Children results.

Similar issues to the non-regularized CFPU reconstruction of the Dancing Children model occur for the SPR method, as shown in the left display of Figure 9, which was obtained with the default



FIGURE 7. Surface reconstructions for the CFPU (top row) and SPR (bottom row) methods for three different surfaces produced from cleaned-up scans of real models (Stanford Dragon, Happy Buddha, and Armadillo).

parameters. These issues can also be resolved by introducing some smoothing, as shown in the right display of this same figure. This latter result was obtained by setting the samplesPerNode parameter to 7. As with the CFPU method with global regularization, this results in over smoothing for the whole surface. The localized regularization approach of the CFPU method offers an effective solution to this problem.

4.6. Computational performance. Timing results (wall clock times) for the CFPU method are displayed in Table 4. Here we have reported the timings for solving the linear systems associated with the potential coefficients in each patch and evaluating these potentials on a uniform grid for the extraction of the level surface using MATLAB's isosurface function (which is based on the Marching Cubes algorithm). As discussed at the end of Section 3.2, the CFPU method is amenable to parallel implementations, so we also report timings with an increasing number of processor cores being used. These results were obtained using MATLAB's Parallel Computing Toolbox. Specifically, we used parfor loops over the patches to solve the linear systems and evaluate the potential. All results were obtained using MATLAB R2020b on a MacBook Pro with 2.4 GHz 8-Core Intel Core i9 processor and 32 GB RAM.

We see from the table that increasing the number of cores from 1 to 2 reduces the computational

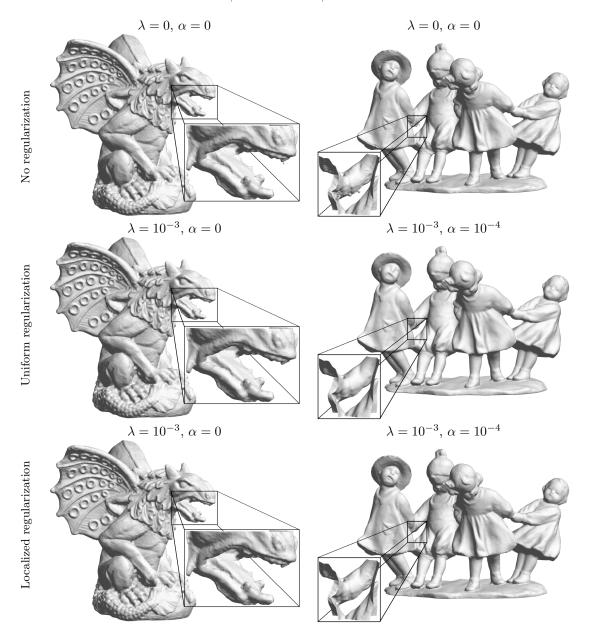


FIGURE 8. CFPU reconstructions of the Gargoyle and Dancing Children models for $\ell=1$. The top row shows the results when no regularization is used, the second row shows the results when the same regularization is used across all patches, and the third row shows the results when using regularization only locally where issues arise in the original reconstructions. The insets of the figures show a more detailed view of the issues.

time considerably, but that these reductions are diminished as the number of cores continues to increase. We believe that this due to non-optimal memory usage associated with parfor and anticipate that a C/C++ implementation using OpenMP will show better scaling. Additionally, we anticipate that the evaluation times (the values in parenthesis) can be further improved by using an octree-based isosurface extraction algorithm, such as [29]. A more detailed study on the computational performance of the method as the number of patches vary with increasing sample points for general

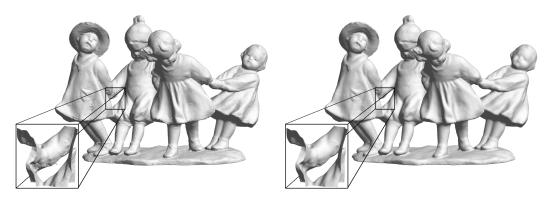


FIGURE 9. SPR results for the Dancing Children model using the default parameters (left) and tuned parameters to resolve the issues around the joined hands.

curl-free RBF interpolants can be found in [17].

5. Discussion and concluding remarks. In this work, we introduced the CFPU method, a novel method that combines curl-free RBFs and a PU framework, for reconstructing an implicit surface from oriented point cloud data. By using PHS kernels for the RBFs, the method avoids the (sometimes) tedious task of selecting good shape or support parameters. However, the method does feature an (integer) smoothness parameter (ℓ). For point clouds with noise and/or sharp features, using $\ell=1$ appears to give good reconstructions of the underlying surface across a range of examples. However, for smooth surfaces without noise in the point clouds, increasing this value appears to result in higher accuracy and convergence rates of the reconstructions. The PU framework allows the method to be computationally efficient for large point sets. Additionally, numerical experiments show that it can represent sharp features of the surfaces.

The CFPU method can provide exact interpolation of a point cloud. Regularization can also be introduced to deal with noise in the normals and the point positions. We found that these regularization methods were effective at eliminating spurious sheets in the zero-level surfaces for point clouds with these issues. We tested the method on a synthetic model and demonstrated its convergence properties. We also tested the method on raw range data and demonstrated the robustness of the method when using PHS of order $\ell=1$ with some regularization. Finally, we presented results for several common 3D models found in the literature. Comparisons of CFPU with the popular SSR and SSD methods show that it can provide similar surface reconstructions of the common 3D models. In the case of the Homer model, which features regions with sparse sampling, CFPU produced smoother reconstructions than SSR and SSD. We additionally found this to be the case for the knot model with synthetically generated noise added to the normals.

The CFPU method is similar to HRBF Implicits [34] and SPR [26,27] in that it uses the oriented normals to determine a potential for the level surface approximating the point cloud. However, there are a few key differences. First, unlike HRBF Implicits, CFPU does not fit the potential directly together with the normal vectors, but exploits properties of matrix-valued curl-free RBFs to extract a potential from a fit of just the normals. It then corrects this potential with a scalar RBF approximant. This two-stage process allows the method to introduce regularizations that can separately deal with noisy normals and noisy point samples. In the case where just the normals are noisy, these regularization can still produce a potential that interpolates the point cloud exactly. Furthermore, the PU framework allows for spatially localized regularization to be applied. Second, CFPU solves for the potential in a strong form instead of a weak form like SPR. This means the normals for the surface do not need to be spread to a 3D background grid and integrated against test functions. The strong form has the added benefit that it can produce surfaces that exactly

interpolate the point cloud. Third, HRBF Implicits and SPR solve large (sparse) linear systems involving all the data in the point cloud to determine an approximating potential. Alternatively, CFPU solves small, decoupled linear systems on each PU patch to determine a local approximating potential that is then blended through a sparse process to determine a global potential. This should make the method more amenable to parallel implementations.

There are several aspects of CFPU that we plan to investigate in the future. One of these is an automated approach for generating the PU patches for non-uniformly sampled point clouds. Another, along these same lines, is to look at ways of adapting the shapes of the PU patches from balls to ellipsoids that better conform to the underlying surface. We also plan to investigate automated approaches for selecting the regularization parameters. While GCV worked well for synthetic noise, we found that it did not produce acceptable results for raw range data. We plan to look into other techniques along the lines of [32] for selecting these parameters. Another regularization approach we will investigate is based on least squares, where fewer RBF centers are used than data sites. This approach has been effective in other RBF-PU methods [31] and has the additional benefit of reducing the computational cost. Finally, we plan to implement a parallel version of the method in C/C++ for even further improvements in computational efficiency.

Acknowledgements. KPD's work was partially supported by a SMART Scholarship, which is funded by The Under Secretary of Defense-Research and Engineering, National Defense Education Program / BA-1, Basic Research. GBW's work was partially supported by National Science Foundation grants 1717556 and 1952674. The Stanford Bunny, Happy Buddha, Stanford Dragon, and Armadillo data were obtained from the Stanford University 3D Scanning Repository. The Homer, Raptor, Filigree, Pump Carter, Dancing Children, and Gargoyle data were obtained from the AIM@SHAPE-VISIONAIR Shape Repository.

REFERENCES

- [1] L. Amodei and M. N. Benbourhim, A vector spline approximation, J. Approx. Theory, 67 (1991), pp. 51–79.
- [2] I. BABUSKA AND J. M. MELENK, The partition of unity method, Int. J. Numer. Meths. Eng., 40 (1997), pp. 727–
- [3] D. M. BATES, M. J. LINDSTROM, G. WAHBA, AND B. S. YANDELL, GCVPACK routines for generalized cross validation, Commun. Statist.-Simula., 16 (1987), pp. 263-297.
- [4] M. N. Benbourhim and A. Bouhamidi, Meshless pseudo-polyharmonic divergence-free and curl-free vector fields approximation, SIAM J. Math. Anal., 42 (2010), pp. 1218–1245.
- [5] M. BERGER, J. A. LEVINE, L. G. NONATO, G. TAUBIN, AND C. T. SILVA, A benchmark for surface reconstruction, ACM Trans. Graph., 32 (2013), pp. 20:1–20:17.
- [6] M. Berger, A. Tagliasacchi, L. M. Seversky, P. Alliez, G. Guennebaud, J. A. Levine, A. Sharf, and C. T. Silva, A survey of surface reconstruction from point clouds, Comput. Graph. Forum, 36 (2017), pp. 301–329.
- [7] F. CALAKLI AND G. TAUBIN, SSD: Smooth signed distance surface reconstruction, Comput. Graph. Forum, 30 (2011), pp. 493-501.
- [8] R. E. CARLSON AND T. A. FOLEY, The parameter r² in multiquadric interpolation, Comput. Math. Appl., 21 (1991), pp. 29-42.
- [9] J. C. CARR, R. K. BEATSON, J. B. CHERRIE, T. J. MITCHELL, W. R. FRIGHT, B. C. McCALLUM, AND T. R. EVANS, Reconstruction and representation of 3D objects with radial basis functions, in Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, 2001, pp. 67–76.
- [10] J. C. CARR, W. R. FRIGHT, AND R. K. BEATSON, Surface interpolation with radial basis functions for medical imaging, IEEE Transactions on Medical Imaging, 16 (1997), pp. 96–107.
- [11] R. CAVORETTO AND A. DE ROSSI, Fast and accurate interpolation of large scattered data sets on the sphere, Comput. Appl. Math., (2010), pp. 1505–1521.
- [12] R. CAVORETTO, A. DE ROSSI, AND E. PERRACCHIONE, Partition of unity interpolation on multivariate convex domains, Int. J. Model. Simul. Sci. Comp., 6 (2015), p. 1550034.
- [13] P. CIGNONI, M. CALLIERI, M. CORSINI, M. DELLEPIANE, F. GANOVELLI, AND G. RANZUGLIA, *MeshLab: an Open-Source Mesh Processing Tool*, in Eurographics Italian Chapter Conference, V. Scarano, R. D. Chiara, and U. Erra, eds., The Eurographics Association, 2008.
- [14] M. CORSINI, P. CIGNONI, AND R. SCOPIGNO, Efficient and flexible sampling with blue noise properties of trian-

- gular meshes, IEEE Trans. Vis. Comput. Graph., 18 (2012), pp. 914-924.
- [15] F. Dodu and C. Rabut, Vectorial interpolation using radial-basis-like functions, Comput. Math. Appl., 43 (2002), pp. 393–411. Radial basis functions and partial differential equations.
- [16] F. Dodu and C. Rabut, Irrotational or divergence-free interpolation, Numer. Math., 98 (2004), pp. 477-498.
- [17] K. P. DRAKE, E. J. FUSELIER, AND G. B. WRIGHT, A partition of unity method for divergence-free or curl-free radial basis function approximation, SIAM J. Sci. Comput., 43 (2021), pp. A1950–A1974.
- [18] J. DUCHON, Splines minimizing rotation-invariant semi-norms in Sobolev space, in Constructive Theory of Functions of Several Variables. Lecture Notes in Mathematics (vol 571), W. Schempp and Z. K., eds., Springer, Berlin, 1977, pp. 85–100.
- [19] G. E. FASSHAUER, Meshfree Approximation Methods with MATLAB, Interdisciplinary Mathematical Sciences, World Scientific Publishers, Singapore, 2007.
- [20] B. Fornberg and N. Flyer, A Primer on Radial Basis Functions with Applications to the Geosciences, SIAM, Philadelphia, 2014.
- [21] E. J. Fuselier, Sobolev-type approximation rates for divergence-free and curl-free RBF interpolants., Math. Comput., 77 (2008), pp. 1407–1423.
- [22] E. J. Fuselier and G. B. Wright, A radial basis function method for computing Helmholtz-Hodge decompositions, IMA J. Numer. Anal., 37 (2017), pp. 774–797.
- [23] J. Gois, V. Polizelli-Junior, T. Etiene, E. Tejada, A. Castelo, L. Nonato, and T. Ertl, Twofold adaptive partition of unity implicits, Vis. Comput., 24 (2008), pp. 1013–1023.
- [24] H. HOPPE, T. DEROSE, T. DUCHAMP, J. McDonald, and W. Stuetzle, Surface reconstruction from unorganized points, SIGGRAPH Comput. Graph., 26 (1992), pp. 71–78.
- [25] S. Jin, R. R. Lewis, and D. West, A comparison of algorithms for vertex normal computation, Vis. Comput., 21 (2005), pp. 71–82.
- [26] M. KAZHDAN, M. BOLITHO, AND H. HOPPE, Poisson surface reconstruction, in Proceedings of the Fourth Eurographics Symposium on Geometry Processing, 2006, pp. 61–70.
- [27] M. KAZHDAN AND H. HOPPE, Screened Poisson surface reconstruction, ACM Trans. Graph., 32 (2013), pp. 1–18.
- [28] M. KAZHDAN AND H. HOPPE, An adaptive multigrid solver for applications in computer graphics, Computer Graphics Forum, 38 (2019), pp. 138–150.
- [29] M. KAZHDAN, A. KLEIN, K. DALAL, AND H. HOPPE, Unconstrained isosurface extraction on arbitrary octrees, in Proceedings of the Fifth Eurographics Symposium on Geometry Processing, 2007, pp. 125–133.
- [30] K. KLASING, D. ALTHOFF, D. WOLLHERR, AND M. BUSS, Comparison of surface normal estimation methods for range sensing applications, in IEEE International Conference on Robotics and Automation, 2009, pp. 3206–3211.
- [31] 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.
- [32] T. C. Lee, Smoothing parameter selection for smoothing splines: a simulation study, Comput. Statist. Data Anal., 42 (2003), pp. 139–148.
- [33] 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.
- [34] I. Macêdo, J. P. Gois, and L. Velho, Hermite radial basis functions implicits, Comput. Graph. Forum, 30 (2011), pp. 27–42.
- [35] B. S. Morse, T. S. Yoo, P. Rheingans, D. T. Chen, and K. R. Subramanian, Interpolating implicit surfaces from scattered surface data using compactly supported radial basis functions, in Proceedings International Conference on Shape Modeling and Applications, 2001, pp. 89–98.
- [36] Y. OHTAKE, A. BELYAEV, M. ALEXA, G. TURK, AND H. P. SEIDEL, Multi-level partition of unity implicits, ACM Trans. Graph., 22 (2003), pp. 463–470.
- [37] Y. Ohtake, A. Belyaev, and H. P. Seidel, 3D scattered data interpolation and approximation with multilevel compactly supported RBFs, Graph. Models, 67 (2005), pp. 150–165.
- [38] 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.
- [39] D. OUYANG AND H.-Y. FENG, On the normal vector estimation for point cloud data from smooth surfaces, Comput. Aided Des., 37 (2005), pp. 1071–1079.
- [40] S. RIPPA, An algorithm for selecting a good value for the parameter c in radial basis function interpolation, Adv. Comp. Math., 11 (1999), pp. 193–210.
- [41] 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.
- [42] V. Shankar and G. B. Wright, Mesh-free semi-Lagrangian methods for transport on a sphere using radial basis functions, J. Comput. Phys., 366 (2018), pp. 170–190.
- [43] D. SHEPARD, A two-dimensional interpolation function for irregularly-spaced data, in Proceedings of the 1968 23rd ACM national conference, New York, NY, 1968, ACM Press, pp. 517–524.
- [44] J. Sussmuth, Q. Meyer, and G. Greiner, Surface reconstruction based on hierarchical floating radial basis

- functions, Comput. Graph. Forum, 29 (2010), pp. 1854-1864.
- [45] I. TOBOR, P. REUTER, AND C. SCHLICK, Reconstructing multi-scale variational partition of unity implicit surfaces with attributes, Graph. Models, 68 (2006), pp. 25–41.
- [46] G. Turk and J. F. O'Brien, Shape transformation using variational implicit functions, in Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '99, 1999, pp. 335–342.
- [47] G. Turk and J. F. O'Brien, Modeling with implicit surfaces that interpolate, ACM Trans. Graph., 21 (2002), pp. 855–873.
- [48] G. WAHBA, Spline models for observational data, CBMS-NSF Regional Conference Series in Applied Mathematics 59, Society for Industrial and Applied Mathematics, 1990.
- [49] C. WALDER, B. SCHÖLKOPF, AND O. CHAPELLE, Implicit surface modeling with a globally regularised basis of compact support, Comput. Graph. Forum, 25 (2006), pp. 635-644.
- [50] H. WENDLAND, Fast evaluation of radial basis functions: Methods based on partition of unity, in Approximation Theory X: Wavelets, Splines, and Applications, Vanderbilt University Press, 2002, pp. 473–483.
- [51] H. WENDLAND, Scattered Data Approximation, Cambridge University Press, Cambridge, 2004.
- [52] G. B. Wright, CFPU Curl-Free RBF Partition of Unity method for implicit surface reconstruction from oriented point clouds, 2022, https://github.com/gradywright/cfpu.
- [53] C. Yuksel, Sample elimination for generating Poisson disk sample sets, Computer Graphics Forum (Proceedings of EUROGRAPHICS 2015), 34 (2015), pp. 25–32.